Addressing Anticipation Effects in Finance

Tomislav Ladika   Elisa Pazaj   Zacharias Sautner

February 6, 2026

Abstract

A wide range of empirical techniques cannot accurately estimate a policy event's causal effects, because agents adjust decisions in advance based on beliefs about future policy outcomes. We show how researchers can measure anticipation bias and refine estimates, by integrating reduced-form and structural estimation. Our novel procedure compares model-predicted outcomes to reduced-form estimates, and only requires a single policy change to implement. We illustrate the importance of this approach by applying it to the Paris Agreement, which is frequently used to understand how agents respond to an increase in climate regulatory risk. We find that before Paris, agents assigned a 77% likelihood to an agreement with some form of emissions penalties being reached. Our estimates imply that anticipation led high-emissions firms to reduce investment and increase cash holdings in the year before the agreement, relative to low-emissions firms. Thus, reduced-form studies of the Paris Agreement may understate its causal effects by up to 50%.

Keywords: Anticipation effects, reduced-form estimation, structural estimation, carbon tax, climate finance

JEL Classification Numbers: G31, C51, Q54.

Ladika is from University of Amsterdam, Email: `t.ladika@uva.nl`; Pazaj is from University of Amsterdam, Email: `e.pazaj@uva.nl`; Sautner (corresponding author) is from University of Zurich and Swiss Finance Institute, Email:`zacharias.sautner@df.uzh.ch`, Postal Address: Plattenstrasse 14, 8032 Zürich, Switzerland.

1 Introduction

The credibility revolution in economics and finance has produced many well-identified empirical estimates of agents' responses to policy events. Yet agents' choices prior to any event depend on their beliefs about possible future policy changes. As such, an estimated response can differ in magnitude and possibly also sign from a policy's causal effect, except in the rare case that the event is completely unexpected (Hennessy and Strebulaev, 2020). Most empirical studies do not correct for such anticipation bias, because agents' beliefs about policy changes usually cannot be measured. Even when such information exists, it usually comes from a different population than the agents who are treated by a specific policy event.

Consider a reduced-form study that examines the impact of electric vehicle (EV) tax credits, by measuring how carmakers' investments change around the announcement of new credits. This approach underestimates the program's causal effect when carmakers anticipated its adoption, and began to invest into EV production before the announcement. Even the estimated sign of the tax effect could be wrong, if carmakers expected larger tax credits than were actually adopted, and subsequently shelved investment after the announcement.

Anticipation effects can also impede the estimation of a structural model of tax credits and EV investment. Uncertainty about the size of future tax credits could depress investment before the program's announcement. If the model omits firms' expectations about tax credits, its estimation may produce higher investment adjustment costs for the pre-announcement period than the true parameter values, in order to match the investment levels observed in the data. In either case, it is challenging for researchers to gauge carmakers' expectations from past policy events, in part because most EV tax credit programs have been adopted very recently.

Hennessy and Strebulaev (2020) offer an important first step towards resolving this challenge. They derive analytical expressions for anticipation bias, and also establish conditions under which an estimate has the same sign as the causal effect. Hennessy and Livdan (2021) extend this analysis to a setting in which agents update their beliefs. Both papers use their formulas to re-examine the impact of corporate taxes on investment. A crucial step is to calibrate firms' beliefs about possible future tax regimes, for which the studies use a long historical time-series of U.S. tax rates. However, this approach cannot be applied to policy events that are historically unprecedented, yet also anticipated in advance. This leaves a significant unaddressed obstacle for much of the current research frontier, which focuses on understanding the impact of such events.

Take the nascent literature examining how investors or firms respond to changes in climate- or nature-related regulations (e.g. Bolton and Kacperczyk, 2021; Garel et al., 2024). Policymakers have contemplated various forms of these regulations, and the possible outcomes have often been reported by financial media. Thus, agents have sufficient information to form beliefs about the range of outcomes before an official policy is announced. Yet these beliefs cannot be calibrated from past policy events, due to the regulations' novelty. This challenge also affects the new literature exploring how artificial intelligence may reshape finance (e.g. Eisfeldt et al., 2025; Babina et al., 2024; Eisfeldt and Schubert, 2024).

This paper provides new guidance on how to credibly estimate causal effects using unprecedented but anticipated policy events, making two primary contributions. First, we show how to account for anticipation bias in such settings, by integrating reduced-form and structural estimation. Our approach uses data on financial outcomes around a policy event, to estimate a distribution of probabilities that agents assigned to possible policy outcomes before the event occurred. It then estimates causal effects accounting for these beliefs.

Second, we apply our new methodology to the Paris Agreement, and study its causal effects on various corporate policies. Negotiated during the COP 21 meeting in 2015, the agreement constitutes the most significant climate policy event to date. Dozens of papers have exploited the event to estimate how changes in environmental regulatory risk affect financial outcomes.[1] We show that firms widely expected some form of climate targets to be agreed upon at COP 21, and as a result standard reduced-form models significantly underestimate the Paris Agreement's causal effects on various firm outcomes.

The general process to implement our integration procedure is as follows. A researcher writes a dynamic model in which the initial policy state can change over time. This can be

[1] Existing work studies changes in the carbon risk premium in equity (e.g., Bolton and Kacperczyk, 2021, 2023) and debt markets (e.g., Degryse et al., 2023; Albuquerque et al., 2025; Seltzer et al., 2025), banks' investments in high-emissions firms (e.g., Alessi et al., 2024), corporate green revenues (e.g., Klausmann et al., 2024), leverage of firms with high climate risk exposures (e.g., Ginglinger and Moreau, 2023), and total factor productivity (e.g., Pang et al., 2023).

a simple extension of an established model, in which some parameter values change with the policy state. At each point in time agents assign probabilities to each possible transition to another policy state (or no state change), modeled using a Markov chain.

Next, the researcher estimates the model using the Simulated Method of Moments (SMM). She conducts numerous simulations, and in each one she solves the model across all time periods. This yields a simulated panel dataset, containing optimal outcomes from both time periods before and after the policy state change. The researcher then estimates a reduced-form model using the simulated dataset, and separately using real data around an observed policy event. We focus on a differences-in-differences (DiD) regression, but event studies or other models can also be used.

The researcher repeats this simulation process for many combinations of transition probabilities and other unknown model parameters, until the simulated DiD coefficients (and other moments) match the empirical data as closely as possible. The final estimates of the unknown parameters are those that produce the best match between the model and agents' observed choices around a policy event. The researcher then quantifies the event's causal effects by re-solving the model in the counterfactual absence of anticipation, with transition probabilities set to 0 and other parameter values set to the final estimates. This is analogous to reduced-form analysis of a natural experiment in which the policy change is exogenous and fully unexpected.

The key novelty is the use of well-identified DiD coefficients as moments to match in SMM. These coefficients provide highly useful information for identifying the transition probabilities and state-contingent parameters. This integration process produces parameter estimates that match the most credible reduced-form evidence available. Understanding the degree of anticipation can be interesting in its own right—for example, it can help policymakers to gauge the magnitude of firms' responses to new regulation. Moreover, counterfactual analysis requires model parameter values that are estimated *in the presence of anticipation*.

We now describe the application of this approach to the Paris Agreement. We develop a simple extension of the standard neoclassical model of corporate investment. We start by modeling a representative firm and introduce two features. First, the firm's profits decrease linearly in a cost parameter representing a carbon tax, which takes a different value in each policy state and is proportional to the firm's capital stock.[2] Second, the firm assigns a probability to each possible transition from the initial policy state to another policy state over the next time increment.[3] The firm continuously chooses optimal investment, given its distribution of beliefs about the future policy state as well as the parameter values in each state. The initial policy state can represent the period before the Paris Agreement when few emissions penalties were in place, but firms anticipated various outcomes of COP 21 and the carbon tax rate that the government would have to adopt given each outcome. A policy state change can represent agreement at COP 21 on a climate goal, which is then implemented via a specific carbon tax.

The baseline model does not capture the full breadth of the Paris Agreement's effects, but it produces closed-form expressions that build intuition about the model's key mechanisms. Optimal investment decreases in the level of the carbon tax, and both investment and firm value are lower than in a model without this parameter—even in the initial state with no carbon tax. The baseline model also shows that the firm faces a transition risk premium, due to the possibility of a value decline when a policy state change leads to a carbon tax. Further, the standard risk premium for cashflow volatility increases with the carbon tax rate.

Environmental policy changes can impact not only a firm's profitability, but also its cost of capital. A key debate in the literature is which of these channels matters more for the sensitivity of firm value to a carbon tax. Thus, we extend the baseline model to include both high- and low-emissions firms, and to incorporate state-contingent financing costs following Bolton et al. (2013) (henceforth "BCW"). Only the high-emissions firm needs to pay a carbon tax following certain policy state changes. As in BCW, fixed and variable costs of financing depend on the policy state and vary for both firms, which in turn choose both investment and cash holdings in anticipation of these changes. Each firm's cost of capital adjusts to the resulting competitive advantage for low-emissions firms.[4] This setup can rep-

[2]A wide range of countries have recently introduced some form of carbon emissions penalty to meet their obligations under the Paris Agreement. Nevertheless, less than 25% of global emissions are currently covered by a carbon tax or cap-and-trade scheme (Pedersen, 2024). Our assumption that the carbon tax increases with the capital stock reflects that emissions typically rise with firm size.

[3]A transition can occur at most once, but the model can easily accommodate multiple state changes. For example, firms can initially hold beliefs about the future adoption of a climate target, and after an agreement's announcement they can form new beliefs about the scale-up or reversal of the policy. Thus, our framework can account for various policy dynamics discussed by Hennessy and Strebulaev (2020).

[4]The importance of financial constraints for understanding the effects of carbon pricing is documented,

resent a government that implements the Paris Agreement by levying a carbon tax on brown firms, and a banking sector that responds by increasing loan markups for these firms.

Because external financing is costly in this full model, both firms may benefit from cutting investment in the current period to increase precautionary savings. Thus, optimal investment now depends on both real and financial frictions, and firm valuations depend on cash holdings in addition to the capital stock. Overall, the model contributes to existing theories of carbon emissions regulation (e.g., Bustamante and Zucchi, 2024; Albuquerque et al., 2025), by accounting for the impact of anticipation on firm investment and financing decisions.

We proceed to estimate the full model using SMM. During simulation, the model's initial *Pre* state (with no carbon taxes) can switch to either a *Mild* or *Strict* state, in which high-emissions firms face a moderate or high carbon tax rate, respectively. Tax rates are calibrated using Social Cost of Carbon (SCC) estimates from Nordhaus (2019), with the *Strict* state corresponding to a goal to limit temperature increases to 2℃. Since our model allows for only one state change, carbon taxes cannot be subsequently increased or watered down.

During the integration step, we estimate DiD regressions following Bolton and Kacperczyk (2021) (henceforth BK) and using both model-simulated and empirical data. BK examines how stock returns change for high- versus low-emissions firms, in the year before versus after the Paris Agreement. We also examine investment and cash holdings as additional outcome variables. Our simulated DiD coefficients closely match those obtained using real data, resulting in highly precise estimates of the two transition probabilities (from the *Pre* to *Mild* and *Pre* to *Strict* policy states). Our process also produces a reasonable match for 14 additional moments, allowing us to estimate other unknown parameters such as the financing costs in each state.

Our results indicate that prior to the Paris Agreement, there was significant anticipation that a carbon tax would be adopted, and that it would most likely be moderate in size. We estimate that firms assigned a 51% probability to the *Mild* state arising, and a 26% probability to the *Strict* state (implying only a 23% likelihood that no carbon tax would be adopted). Thus, a relatively high level of anticipation is necessary to reconcile our

for example, in Döttling and Rola-Janicka (2025).

model with firms' observed responses to the Paris Agreement.[5] Interestingly, despite this anticipation, investment rates of high-emissions firms declined significantly after the Paris Agreement relative to those of low-emissions firms. This indicates that the agreement is consistent with the *Strict* state, since less ambitious climate targets would have been viewed as better-than-anticipated news by high-emissions firms (and worse-than-anticipated news for low-emissions firms), and their relative investment rates would have increased.

Our counterfactual exercise solves the model for a completely unexpected switch from the *Pre* to *Strict* state, and compares the resulting DiD regression results to those obtained from model estimation. The decrease in high-emissions firms' investment would be 35% larger following an unanticipated increase in carbon taxes, while the rise in cash holdings would be 50% smaller (both relative to low-emissions firms). This indicates that anticipation of a regulatory change led brown firms to already reduce investment, and engage in precautionary savings, prior to the Paris Agreement. Without accounting for such changes, standard DiD models likely underestimate the agreement's causal effects on these outcomes.

We contribute to a nascent literature showing how empirical research can benefit from the integration of reduced-form and structural estimation (Whited, 2023). This process can help researchers to understand the determinants of an estimated reduced-form elasticity (Briggs et al., 2021) or to rule out a violation of the compound exclusion restriction when re-using a natural experiment (Cronqvist et al., 2024). Our paper shows how integration can also help researchers to account for anticipation bias in many settings. Some work attempts to estimate transition probabilities by combining stock and option prices (e.g., Subramanian, 2004; Borochin and Golec, 2016). This approach requires data on multiple near-the-money option contracts, and relies only on market prices from a few days around a policy event. Because our approach does not face these limitations, it can potentially estimate more unknown parameters and be applied to a wider range of empirical settings. More broadly, we contribute to a growing literature on improving the credibility of reduced-form estimates in finance (Baker et al., 2022; Heath et al., 2023).

Our carbon tax application contributes to the climate finance literature, in particular

[5]Consistent with such a high level of anticipation, U.S. Secretary of State John Kerry stated just six weeks before the Paris meeting that diplomats expected "the most significant international agreement on the issue ever reached" (Worland, 2015).

studies that examine how financial outcomes change around environmental regulations. Our findings caution that some results may need to be reinterpreted when anticipation is not properly accounted for, if the goal is to understand whether and how climate regulatory risks affect firms and financial markets (and to quantify such effects). Existing work proposes a "green paradox" that high-emissions firms may increase investment in anticipation of climate regulations (e.g., Albuquerque et al., 2025), and we complement it by showing how to incorporate anticipation into both theoretical modeling and empirical estimation.

2 Model

2.1 Overview of modeling approach

We present a framework that adds a carbon emissions penalty to a standard neoclassical model of investment. Time is continuous and denoted by $t \geq 0$, with each time increment denoted by dt. At $t = 0$ no emissions penalties are in place. Future penalty levels are uncertain and vary by regulatory policy state, and market participants possess expectations about the likelihood of a transition to another regulatory policy state. The model's purpose is to study how carbon penalties affect firm investment in the presence of such anticipation effects.

We first develop a benchmark case with a single representative firm that faces no financing frictions. This setup yields analytical expressions for the firm's first-best investment policy and its elasticity to carbon taxes. We then present the full model that includes two firms with different emissions levels, and also contains external financing costs that vary with the firm's type and the policy state. This full model provides economic structure for a DiD comparison of high- and low-emission firms around significant climate policy changes, such as the Paris Agreement.

2.2 Benchmark: Neoclassical model with uncertain carbon taxes

2.2.1 Benchmark model setup

Carbon emissions penalty. The emissions penalty is $\tau_s^e K_t$, where τ_s^e is a tax on carbon emissions, K_t the firm's capital stock at time t, and s the climate regulatory policy state.

The penalty represents that carbon taxes can change over time and increase with the capital intensity of business activities. The regulatory policy state is captured by a climate state variable s_t, which follows a time-homogeneous Markov chain. There are N possible regulatory states. The economy begins in State 1 with no carbon taxes ($\tau_1^e = 0$). This can represent an initial state before any major climate regulations are adopted, and is analogous to the *Pre* period of a DiD model analyzing the Paris Agreement. In the other $N-1$ states carbon taxes are positive, with $0 = \tau_1^e < \tau_2^e < \cdots < \tau_N^e$.

Given an economy in regulatory policy state s at time t, firms and investors assign probabilities of an instantaneous transition to each of the other states s^-, denoted by $\zeta_{s,s^-} \in [0,1)$. Because ζ_{s,s^-} represents the likelihood of an instantaneous change, its inverse equals the duration of the current policy state s. Policy state changes are intended to be infrequent, but can lead to potentially large changes in carbon taxes. As we show below, firms take the likelihood of transitioning out of the current state into account when making investment decisions.[6]

Dynamics of capital and profits. At each time t the firm invests into new capital, earns profits, and pays out dividends. The dynamics of the firm's capital stock are standard and given by:

$$dK_t = (I_t - \delta K_t)\,dt, \tag{1}$$

where δ denotes the rate of capital depreciation and I_t total investment.

The firm's incremental operating profit dY_t also has a standard form, except that it decreases with the level of carbon taxes:

$$dY_t = dA_t K_t - I_t dt - \Gamma(I_t, K_t)dt - \tau_s^e K_t dt, \tag{2}$$

where dA_t is a stochastic shock to firm productivity that we define below, and $\Gamma(I_t, K_t)$

[6] Hennessy and Strebulaev (2020) derive a linear system for the marginal values of adjusting a firm outcome variable in each of the N possible policy states (e.g., the marginal value of increasing investment under each possible tax regime). The system depends on an $N \times N$ "augmented transition matrix", in which each cell contains i) the probability that a shock will occur, if the current policy state is i; and ii) the conditional probability that the policy state will change from i to j, given the shock. The causal effect depends on only the change in benefits as policy changes from one state to another. However, the observed change in outcome variable depends on both the agents' precedent beliefs that the shock would occur, and their subsequent beliefs about further policy shocks. This approach cannot be used when the transition intensity matrix cannot be estimated outside the model (which is the case for unprecedented events, for which there is no historical data).

is a standard investment adjustment cost function. This function is homogeneous of degree one in capital and investment (as in Hayashi, 1982), and thus can be expressed as $\Gamma(I_t, K_t) = g(i_t)K_t$, where $i_t = \frac{I_t}{K_t}$ is the investment to capital ratio. The function $g(i_t)$ is increasing and convex in i_t, and it has a quadratic form:

$$g(i_t) = \frac{\theta}{2}(i_t - \nu)^2, \tag{3}$$

where $\theta > 0$ measures the degree of the adjustment cost and ν is a centering parameter for adjustment costs. The formulation in Eq. (3) is often used in the literature to represent a firm that incurs adjustment costs only when increasing its capital stock, rather than replacing depreciated assets.

In the benchmark model, all operating profits net of investment costs and carbon taxes are paid out as dividends, with dD_t denoting the incremental payout to shareholders. The firm also chooses to optimally liquidate at any time τ, in which case its assets recover value $L_\tau = \omega K_\tau$ with $\omega \in [0, 1]$. This value is constant across regulatory policy states, because the model's focus is on studying how firm outcomes vary with carbon emission penalties (and external financing costs in Section 2.3).[7]

Firm value. The firm's value equals the expectation of all discounted proceeds to shareholders until liquidation:

$$V(K_t) = \mathbb{E}^{\mathbb{Q}}\left[\int_0^\tau e^{\int_0^t r_u d_u} dD_t + e^{\int_0^\tau r_u d_u}\omega K_\tau\right], \tag{4}$$

where the expectation is taken under the risk-neutral measure $\mathbb{Q}$, and $e^{\int_0^t r_u d_u}$ represents the cumulative discount factor for firm payouts.

We assume that investors are risk averse and that there is a constant market price of risk η. The risk-adjusted productivity shock under the risk-neutral measure $\mathbb{Q}$ is:

$$dA_t = \hat{\mu}dt + \sigma d\hat{W}_t, \tag{5}$$

where $\hat{W}_t$ is a standard Brownian motion under the risk neutral measure $\mathbb{Q}$ and $\hat{\mu}$ is the

[7]This assumption is straightforward to relax, to allow for the possibility that assets which produce high emissions have lower liquidation values in a regulatory policy state with emissions penalties.

risk-adjusted mean productivity shock. The following risk-adjustment applies:

$$\hat{\mu} = \mu - \rho\sigma\eta, \tag{6}$$

where ρ is the correlation between firm and aggregate cashflows, and η the market price of risk.

2.2.2 Benchmark model problem

The firm chooses its investment policy to maximize firm value from Eq. (4), with net operating profits continuously paid out as dividends dD_t. If at time $t = \tau$ the firm experiences a sufficiently negatively productivity shock such that $V(K_\tau)$ falls below the liquidation value L_τ, then the firm optimally chooses to liquidate.

Using the principles of dynamic programming, the firm's problem can be re-expressed as the following system of Hamilton-Jacobi-Bellman (HJB) equations:

$$rV(K,s) = \max_{i_s} \left\{ \hat{\mu}K - i_s K - \frac{\theta i_s^2}{2} - \tau_s^e K + V_K(K,s)(i_s - \delta)K + \sum_{s=1}^{N-1} \zeta_{s,s^-} \left[V(K,s^-) - V(K,s)\right] \right\} \tag{7}$$

where the time subscript t is omitted for notational convenience.

The left-hand side (LHS) of Eq. (7) represents the firm's required rate of return, which under the risk-neutral measure $\mathbb{Q}$ is equal to the risk-free rate r. The right-hand side (RHS) shows that firm value is increasing in incremental profits, which depend on the cashflow growth rate $\hat{\mu}K$, net investment $(i - \frac{\theta i^2}{2})K$, and the carbon tax $\tau_s^e K$. The term $V_K(K,s)(i - \delta)K$ captures the effect on firm value of changes in the capital stock due to investment and depreciation. The final term in Eq. (7) captures the effect of anticipation of a possible change in the climate regulatory state. The effect depends on the likelihood of a transition to another state and the sensitivity of firm value to such a transition, captured by $V(K,s^-) - V(K,s)$.

2.2.3 Benchmark model solution

The first-order condition on the above HJB yields:

$$q_s = 1 + \theta i_s, \tag{8}$$

where $q_s = V_K(K, s)$ is the firm's marginal benefit from creating an additional unit of capital in state s. The RHS of Eq. (8) is the marginal cost of creating one unit of capital, which equals the replacement cost (normalized to 1) plus the firm's investment ratio i_s scaled by its adjustment cost parameter θ. This expression shows that in each state s, the firm optimally sets i_s so that the marginal benefit of capital equals its marginal cost.

The first-best investment policy is given by:

$$\begin{aligned} i_s^{FB} = &\left(r + \delta + \sum_{s \neq s^-} \zeta_{s,s^-}\right) \\ &- \sqrt{\left(r + \delta + \sum_{s \neq s^-} \zeta_{s,s^-}\right)^2 - \frac{2\{\hat{\mu} - [\tau_s^e + r + \delta + \sum_{s=1}^{N-1} \zeta_{s,s^-}(1 - q_{s^-})]\}}{\theta}} \end{aligned} \tag{9}$$

where q_{s^-} denotes marginal q in each other state s^-. Average q and marginal q are equal due to the homogeneity property of adjustment costs, so q_{s^-} also represents the scaled firm value function $V(K, s^-)/K$.

Eq. (9) implies that gross investment i_s^{FB} is positive only if $\hat{\mu} > \tau_s^e + r + \delta + \sum_{s \neq s^-} \zeta_{s,s^-}(1 - q_{s^-})$. When there is no uncertainty regarding the climate regulatory state ($\zeta_{s,s^-} = 0 \, \forall \, s$), the return on investment needs to exceed the neoclassical benchmark of $(r + \delta)$ by the amount of the carbon tax τ_s^e. The effect of uncertainty on investment depends on the likelihood of each regulatory state and on the marginal q_{s^-} values in each state. When the transition-intensity weighted average of the q_{s^-} values is less than 1, then the term $\sum_{s=1}^{N-1} \zeta_{s,s^-}(1 - q_{s^-})$ is positive, and the threshold for a positive investment policy increases further.

The elasticity of investment with respect to the carbon tax can be derived from Eq. (9) as:

$$\frac{\partial\, i^{FB}}{\partial\, \tau_S^e} = -\frac{1}{\theta\sqrt{\left(r + \delta + \sum_{s \neq s^-} \zeta_{s,s^-}\right)^2 - \frac{2\{\hat{\mu} - [\tau^e + r + \delta + \sum_{s=1}^{N-1} \zeta_{s,s^-}(1 - q_{s^-})]\}}{\theta}}} \tag{10}$$

This expression represents the causal effect of an tax policy on investment in the presence of anticipation, and it is always negative since $\theta > 0$. This means that an increase in the emissions tax level leads to a reduction in the first-best investment policy. Optimal investment and marginal (and average) q are therefore lower compared to the neoclassical benchmark (where $\tau_s^e = 0\,\forall s$ and there is no uncertainty regarding the climate regulation state).

Interestingly, our results differ from Albuquerque et al. (2025), who find that high-emissions firms may optimally *increase* investment prior to adoption of a carbon tax. One reason is that the carbon tax in our model is proportional to total capital, so pre-tax investment reduces long-term profitability, while the emissions penalty in their model is a constant. Another is that a policy state change can occur at any time in our model, compared to just a single interim period in their setup.

The anticipation of a regulatory state change also impacts the firm's risk premium $\mu_R(s) - r$, expressed as:

$$\mu_R(s) - r = \frac{\rho\sigma\eta}{q(s)} - \sum_{s=1}^{N-1} \frac{(e^{\kappa_s} - 1)\zeta_{s,s^-}(q(s^-) - q(s))}{q(s)}, \tag{11}$$

where $\mu_R(s)$ denotes the firm's required rate of return in state s, and κ_s captures the risk-adjustment for the transition intensity out of state s, such that $\hat{\zeta}_{s,s^-} \equiv e^{\kappa_s}\zeta_{s,s^-}$ denotes the transition intensity under the risk-neutral measure.

Eq. (11) shows that the firm's total risk premium consists of two components. The first term on the RHS is a productivity risk premium, which is increasing in the correlation between firm and aggregate cashflows ρ, the volatility of firm cashflows σ, and the market price of productivity risk η. It is also inversely related to marginal q, implying that return premia are lower for high-investment firms (consistent with Hou et al., 2015; Fama and French, 2015). Moreover because q and carbon taxes are inversely related, a higher carbon tax level also leads to a higher productivity risk premium.

The second term on the RHS of Eq. (11) is a premium for uncertainty about the future climate regulatory state, often referred to as a "transition risk premium" in the climate finance literature. To understand the intuition for this premium, first consider a case in which firm value is higher in state s than other states s^-, i.e., climate regulation has a negative effect on firm value. Then $\kappa_s > 0$ and the second term is positive overall (since $e^{\kappa_s} - 1 > 0$ and $q(s^-) - q(s) < 0$). This implies that anticipation of potential climate regulation in the future increases the transition risk premium in the current state.[8] The opposite is true when firm value is lower in state s than other states, i.e., climate regulation boosts firm value. The term is negative overall, implying that anticipation of a potential switch to better states decreases the transition risk premium. Uncertainty about the climate regulation state therefore leads to greater return volatility, consistent with Pástor and Veronesi (2013).

2.3 Full model with financing frictions

2.3.1 Full model setup

We now extend the benchmark model to two firms and allow for external financing. Firms make investment and financing decisions jointly (similar to Bolton et al., 2011), and we introduce uncertainty about financing costs in a similar spirit to Bolton et al. (2013) to explore how climate regulations affect financing conditions. This full model better captures reality, and thus its estimation can provide more precise parameter estimates. This is important for quantifying the causal effect of a climate regulation, as well as market participants' prior expectations that the change would occur.

Firm types. There are two types of firms, green (G) and brown (B), that differ based on the emissions intensity of their capital. The emissions penalty is now expressed as $\tau^e_{f,s} K_t$, where $\tau^e_{f,s}$ is the carbon tax for firm $f \in \{G, B\}$. We normalize $\tau^e_{G,s}$ to 0 in all policy states, while $\tau^e_{B,s}$ may be positive in some states. This differentiation between green and brown firms loosely follows Pedersen (2024), whereby green capital does not produce any emissions.

Financing costs. Firms can raise costly external equity financing at any time t. Following Bolton et al. (2013), we assume a linear cost structure for equity issuance. In each state s, there is a fixed financing cost of $\phi_{f,s}$ and a marginal cost of $\gamma_{f,s}$. State-contingent financing costs represent the possibility that in a regulatory state with high emissions penalties, brown firms face more difficulty raising equity financing (e.g., due to higher information

[8] It also implies that the risk-neutral transition intensity $\hat{\zeta}_{s,s^-}$ is greater than the intensity under the physical measure. This intuitively captures the idea that a risk-averse investor perceives states which are good for firm value as being of shorter duration than a risk-neutral investor would.

asymmetry about their future prospects) while green firms may be able to raise equity at lower cost (e.g., due to higher demand from sustainability-focused investment funds). We denote the cumulative external financing process by E, with dE_t denoting financing raised in each time increment.

Cash savings. In some time increments the firm's desired investment spending may exceed its current operating profits, especially given that those profits can be negative. The firm can make up the gap by raising external financing. Yet since this is costly, the firm may instead prefer to use internal cash savings. Thus, in the full model the firm has a precautionary motive to maintain a cash inventory M_t. Cash accumulation has the following dynamics:

$$dM_t = dY_t + (r - \lambda)M_t dt - dD_t + dE_t, \tag{12}$$

where λ is the carry cost of cash. The firm's return on its cash balances $(r - \lambda)$ is lower than the risk-free rate, which can reflect agency costs of hoarding cash.

2.3.2 Full model problem

Unlike in Section 2.2, the firm does not pay out all excess profits as dividends. Instead, in each time increment the firm's operating profits dY_t and external financing raised dE_t equal the sum of its investment spending, dividend payouts dD_t, and cash savings dM_t. The firm chooses an investment I, payout D, and external financing E policy, along with liquidation time τ, to maximize its value:

$$\mathbb{E}^{\mathbb{Q}}\left[\int_0^{\tau} e^{-rt}(dD_t - dE_t - \phi_{f,s}K_t 1_{dE_t>0} - \gamma_{f,s}dE_t) + e^{-r\tau}(\omega K_\tau + M_\tau)\right] \tag{13}$$

where the first term represents the value of discounted net payouts to shareholders, and the second term represents the value of the firm upon liquidation.

Eq. (13) shows that in each state of the world, firm value depends on the cash stock M_t and capital stock K_t. Let $V(M, K, s)$ denote the value of the firm in state s. The firm's optimal policies in each state imply an upper payout boundary $\overline{M}_s$, such that whenever $M_t > \overline{M}_s$ the firms pays out excess cash as a dividend to shareholders. Similarly, there is a lower boundary $\underline{M}_s$ such that the firm liquidates as soon as $M_t < \underline{M}_s$.[9] In the interior region $M \in$

[9]When cash holdings become perilously low, the firm can either liquidate or raise significant external

$(\underline{M}_s, \overline{M}_s)$, firm value satisfies the following HJB equation under the risk-neutral measure:

$$
\begin{aligned}
rV(M,K,s) = \max_{i_s} \quad & \Bigg\{ \Big[\hat{\mu}K - i_s K - \frac{\theta i_s^2}{2} - \tau^e_{f,s} K + (r-\lambda)M \Big] V_M(M,K,s) \\
& + \frac{1}{2} K \, \sigma^2 V_{MM}(M,K,s) + V_K(M,K,s)(i_s - \delta)K \\
& + \sum_{s=1}^{N-1} \hat{\zeta}_{s,s^-} \left[V(M,K,s^-) - V(M,K,s) \right] \Bigg\}
\end{aligned} \tag{14}
$$

Similar to the HJB (7) of the benchmark model, the RHS of Eq. (14) details how each of the state variables (M, K) and the regulatory state s affect firm value. The first term represents the effect of a change in cash holdings M, which depends on the expected cashflow growth rate $\hat{\mu}K$, total investment and associated adjustment costs $i_s K + \theta i_s^2/2$, the carbon tax $\tau^e_{f,s} K$, and the net interest earned on cash balances $(r - \lambda)M$. A higher carbon tax diminishes the marginal effect of cash accumulation on firm value. The second term on the RHS represents the valuation effect of the volatility in cash holdings, while the third term captures the effect of changes in the capital stock due to investment and depreciation.

The final term in Eq. (14) captures anticipation effects in the full model. As in Section 2.2.2, this term depends on the transition intensity $\hat{\zeta}_{s^-,s}$ from the current state s to another state s^-, as well as the expected change in firm value due to the state change, $V(M,K,s^-) - V(M,K,s)$. However unlike in the benchmark model, this sensitivity of firm value to the state change now depends on the firm's cash holdings as well as its physical capital. This can represent that firm value decreases by more following the introduction of carbon taxes for firms that have lower precautionary cash savings, and thus are more exposed to changes in financing costs.

2.3.3 Full model solution

To solve the model, we make use of the homogeneity property of the value function, whereby firm value in each state s is homogeneous of degree one in both cash and capital:

$$
V(M,K,s) = KV(\frac{M}{K}, s) \equiv K \; F(m,s), \tag{15}
$$

financing. The threshold $\underline{M}_s$ arises when liquidation provides greater firm value than continuation after raising costly financing.

where $m \equiv M/K$ is the cash-to-capital ratio and $F(m,s)$ is the scaled value function. This homogeneity property yields an analytically tractable framework that allows us to study the impact of changes in pollution costs on firms' investment decisions, taking into account the anticipation of changes in climate regulation states. Note also that $V_K = F(m,s) - mF'(m,s)$, $V_M = F'(m,s)$, and $V_{MM} = \frac{1}{K}F''(m,s)$.

Given these properties, the HJB can be re-written as:

$$rF(m,s) = \left[\hat{\mu} - i_s - \frac{\theta i_s^2}{2} - \tau_{f,s}^e + (r-\lambda)m\right] F'(m,s) + \frac{1}{2}\sigma^2 F''(m,s) \\ + (i_s - \delta)\left[F(m,s) - mF'(m,s)\right] + \sum_{s=1}^{N-1} \zeta_{s,\hat{s^-}} \left[F(m,s^-) - F(m,s)\right]$$

Re-arranging this expression yields the following second-order coupled ordinary differential equation:

$$(r - i_s - \delta)F(m,s) = \left[\hat{\mu} - \frac{\theta i_s^2}{2} - \tau_{f,s}^e + (r - \lambda - i_s - \delta)m\right] F'(m,s) \\ + \frac{1}{2}\sigma^2 F''(m,s) + \sum_{s=1}^{N-1} \zeta_{s,\hat{s^-}} \left[F(m,s^-) - F(m,s)\right] \tag{16}$$

We derive the firm's first-best choices by solving Eq. (16) numerically, using the upper and lower boundary conditions on cash m for each state s. The first-order condition for investment is given by:

$$i^*(m,s) = \frac{1}{\theta}\left(\frac{F(m,s)}{F'(m,s)} - m - 1\right) \tag{17}$$

Eq. (17) shows that investment, which is a function of scaled cash m and the regulatory state s, depends on both real and financial frictions. As in the benchmark model, investment is decreasing in adjustment costs θ. However now investment also decreases as financing constraints become more binding. To see this, note that the scaled value function $F(m,s)$ is increasing and concave in m, while its first derivative $F'(m,s)$ is decreasing and convex. Therefore, when a firm is liquidity constrained (m is low), its scaled value will be low and the marginal value of cash will be high. This implies a low value of the ratio $F(m,s)/F'(m,s)$, and therefore of optimal investment $i^*(m,s)$.

3 Integration procedure: Overview and benefits

Our approach of integrating DiD and structural estimation adds three features to a standard SMM procedure.[10] First, state changes occur during the model's simulation, to mimic the policy shocks that DiD models typically use for identification. Second, the researcher specifies firms' belief structure about the state changes, and simulates their choices given these beliefs. Third, DiD coefficient estimates are used as moments to help identify the unknown model parameters. This section details these three features and discusses their benefits. It also explains the advantages of our procedure over a calibration approach.

3.1 Specifying transition intensities and simulating state changes

Our general approach begins with an economic model in which some parameter values vary with the policy state, and discrete state changes can occur one or more times. A state change's arrival is modeled with a Markov chain, which is a matrix of transition probabilities of a change from each state s to another state s^- during a single time increment. For example, when the economy can be in one of three possible states (s_1, s_2, s_3), the transition matrix is:

$$\begin{bmatrix} \zeta_{s_1,s_1} & \zeta_{s_1,s_2} & \zeta_{s_1,s_3} \\ \zeta_{s_2,s_1} & \zeta_{s_2,s_2} & \zeta_{s_2,s_3} \\ \zeta_{s_3,s_1} & \zeta_{s_3,s_2} & \zeta_{s_3,s_3} \end{bmatrix}$$

where the element in row i and column j contains the transition intensity from state i to state j. Each row i contains the transition intensities to each possible state (including the probability ζ_{s_i,s_i} of no immediate state change), so its values add up to 1.

Transition intensities are generally not observable, but can be estimated using SMM. Let ψ denote the set of all unknown model parameters to estimate, and $\hat{\psi}$ the set of values for these parameters during the current SMM iteration. In each simulation, a state change can arise in each time increment based on a stochastic process, which depends on the transition

[10] Appendix IA.5 provides a general overview of SMM, intended for readers who are unfamiliar with the procedure. Economic models can also be estimated using the Generalized Method of Moments, but we do not use this approach as our model does not produce closed-form expressions for some moments.

intensity values in $\hat{\psi}$. When a state change occurs, it applies to all firms at the same time. The timing of a state change differs in each simulation, and it is possible for no state changes to occur in some simulations.

A researcher can impose structure on the Markov chain to determine the number of transition intensities to be estimated. First, she can choose the number of possible states, which pins down the size of the matrix. Second, she can pre-set some matrix elements by making assumptions about transitions between certain states. For example, setting $\zeta_{s_3,s_3} = 1$ (and thus $\zeta_{s_3,s_1} = \zeta_{s_3,s_2} = 0$) implies that the economy remains in state s_3 permanently once it arises.

3.2 Using DiD coefficients as moments

In our approach, the moments used in the SMM process include coefficient estimates from DiD regressions run on both simulated and empirical data. Thus, the simulated economy should correspond to the DiD model's setup by including multiple firm types.[11] Also, a state change's impact on the parameters in ψ should differ across firm types. For example, in our application to the Paris Agreement, we simulate an economy that contains brown and green firms, and a change from the *Pre* state leads to a carbon tax levied only on brown firms.

During each simulation, a researcher computes optimal outcomes for each firm type, in all time periods during which the economy is in its initial state and also after each state change. This yields a panel dataset of simulated model outcomes. Next, the researcher estimates DiD regressions using this simulated data, and the interaction term coefficient estimate is included in the vector of simulated moments (in some applications, other DiD coefficients could also be used as moments). Separately, the researcher estimates the same DiD regression using actual data, and includes the interaction coefficient estimate in the vector of empirical moments. Given the current values in $\hat{\psi}$, SMM determines the next guess for each parameter by solving an objective function that minimizes the distance between the simulated and empirical moments (see Eq. (IA.3) in Appendix IA.5).[12]

[11] DiD regressions can also sometimes be integrated with structural estimation of a model with a single representative firm. When a policy state change does not affect control firms along any dimension, then changes in a representative firm's outcomes around the policy event can be matched to reduced-form DiD coefficients.

[12] Our integration procedure can be applied to other reduced-form methodologies. For example, a researcher can write a model in which a key outcome is the firm's valuation. Simulation produces a dataset of firm values for each time period, which can be used to calculate simulated stock returns. The researcher

The moments obtained from DiD regressions are highly informative for identifying unknown state-contingent parameters (including the transition intensities), because changes to real firm outcomes after an actual policy shock are significantly influenced by these parameters. DiD moments can also be highly useful for identifying some parameters that do not vary with the policy state. In our application, the sensitivity of investment rates to a carbon tax likely depends on firms' ex-ante beliefs that the Paris Agreement would be adopted, as well as the size of their time-invariant investment adjustment costs.

One important consideration for our procedure is that model simulation can produce state changes which have not occurred in real life. Consider a researcher who uses a three-state Markov chain (as in Section 3.1), and initially guesses non-zero values for ζ_{s_1,s_2} and ζ_{s_1,s_3} in $\hat{\psi}$. In the first set of simulations, the economy's state will change from s_1 to s_2 in some individual simulations, but to s_3 in others. This yields simulated data on firm outcomes in all three states, while the actual data only contains outcomes for one state before and one state after a policy shock. In other words, some simulations produce counterfactual states that real-life market participants have anticipated but not experienced. This situation does not arise when estimating models without anticipation, but that is because those processes overlook the fact that market participants' beliefs affect the outcomes observed in actual data.

Our procedure maps the observed post-shock state to one of the states in the simulation, while designating other states as counterfactuals. For example, a researcher can determine that the actual policy shock corresponds to a change from state s_1 to s_3, and that s_2 was anticipated but did not occur. During each set of simulations, the panel dataset is constructed using only individual simulations with a state change from s_1 to s_3. The DiD coefficient estimates and other simulated moments are constructed using only these data, while data from simulations in which the counterfactual state s_2 arises (or no state change occurs) are discarded.[13] Importantly, the data from retained simulations include firm outcomes that depend on anticipation that the state could switch to s_2. Thus, simulated moments can still be used to identify the transition intensity ζ_{s_1,s_2}.

can then estimate the cumulative return around a state change using an event study framework, and match this to the corresponding estimate from an event study run on actual data.

[13]Section 4.3.1 discusses an empirical test to check whether the realized and counterfactual states are correctly classified.

3.3 Comparison with calibration

Hennessy and Strebulaev (2020) and Hennessy and Livdan (2021) provide important contributions for accounting for the effects of anticipation using calibration rather than estimation. Calibration can be informative for understanding the qualitative implications of a model, while estimation is helpful for learning about unknown model parameters and for generating quantitative implications.[14] In addition, estimation can yield statistics regarding the uncertainty of parameter estimates. Further, it can match the actual data along multiple dimensions (at least as many as the number of unknown model parameters), while calibration typically matches only a few, often ad hoc, moments.

The limitations of applying calibration to our model are evident in the elasticity of investment with respect to the carbon tax rate in Eq. (10), which depends on the values of all model parameters including the unknown transition intensities. Counterfactual exercises therefore heavily rely on these parameter values. Another limitation is that treated firms in most DiD analyses are usually different from the average sample firm, but estimates of economic model parameters are often available only for the latter. As such, existing literature provides little guidance for how to pre-set the values even of standard parameters for green and brown firms. Finally, when the closed-form solutions from Hennessy and Strebulaev (2020) and Hennessy and Livdan (2021) are applied to a complex setting that does not yield analytical expressions (such as the impact of external financing costs around the Paris Agreement), then the causal effects inferred from shock responses may be inaccurate.

4 Economic model estimation

This section applies our procedure of integrating DiD and structural estimation to the full economic model developed in Section 2.3. We first describe the sample and data used to calculate empirical moments. Second, we discuss the estimation process for unknown model parameters, and our choices of pre-set values for other parameters. Third, we explain our

[14]A researcher conducts a calibration by pre-setting the values of *all* fundamental parameters in an economic model. The researcher then studies how the firm's optimal choices change as model parameters are adjusted. While this process shares some similarities with counterfactual analysis in structural estimation, it usually does not produce quantitative estimates of causal effects.

choice of moments to identify these unknown parameters, and present results from DiD regressions that examine changes to various firm outcomes around the Paris Agreement. Fourth, we report the results of the model's estimation, including the transition intensity estimates. Fifth, we conduct a counterfactual analysis to quantify the Paris Agreement's causal effects on firm outcomes, after accounting for anticipation effects.

4.1 Sample

We obtain data on quarterly fundamentals from Compustat, monthly stock returns from CRSP, and annual carbon emissions from S&P Trucost. Our sample contains 457 U.S. firms in the intersection of these databases, and covers the last three months (quarter Q4) of 2014 through Q4 of 2016. We use this sample to construct empirical moments, including the DiD estimates in Section 4.3. We omit Q4 of 2015, since it is not clear whether firm outcomes were determined before or after the Paris Agreement's announcement in the middle of the quarter. Table 1 presents summary statistics for the variables used for empirical moments.

We sort sample firms into a treatment and control group, corresponding to brown and green firms in the model. We first measure each firm's emissions intensity as Scope 1 emissions scaled by total revenue, both measured in 2014 to account for the typical lag in reporting of emissions data identified by Zhang (2025). Scope 1 emissions are the tons of carbon dioxide equivalent that a firm produces directly through its operations. Emissions scaled by revenues are more informative for comparing firms of different sizes than absolute emissions (Zhang, 2025). Next, we classify a firm as treated if its emissions intensity is above the 66^{th} percentile in 2014. Further, we follow BK to construct the control group by matching each treated firm to one firm with an emissions intensity below the 66^{th} percentile, using the nearest neighbor methodology without replacement.[15] We also use an alternative classification which is identical except that it uses the 75^{th} percentile of emissions.

[15] We construct propensity scores by matching on all variables in Table 1 except *Cash* and *Returns*, which are firm outcome variables in our analysis but are not used in the matching process of BK.

4.2 Model Parameters

4.2.1 Parameters to be estimated

The transition intensities, state-contingent financing costs, and investment adjustment costs from our model are difficult to measure outside the model, and there is little evidence in existing literature on their values (especially for the brown and green firms in our setting). Hence, we cannot reliably pre-set these parameters' values. Instead, we include 12 unknown parameters in our set ψ to estimate:

$$\psi = \{ \zeta_{Pre,Mild},\ \zeta_{Pre,Strict},\ \theta_G,\ \theta_B, \gamma_G,\ \phi_{G,Pre},\ \phi_{G,Mild},\ \phi_{G,Strict},\ \gamma_B,\ \phi_{B,Pre},\ \phi_{B,Mild},\ \phi_{B,Strict} \}$$

The Markov chain that governs transitions between the policy states is expressed as a 3×3 matrix, whose rows/columns correspond to the *Pre*, *Mild*, and *Strict* states, respectively:

$$\begin{bmatrix} \zeta_{Pre,Pre} & \zeta_{Pre,Mild} & \zeta_{Pre,Strict} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Under this Markov chain the policy state can only change once. Specifically, a change from the *Pre* state to either the *Mild* or *Strict* state is permanent: $\zeta_{Mild,Mild} = \zeta_{Strict,Strict} = 1$. We further do not contemplate the possibility that emissions penalties are later tightened ($\zeta_{Mild,Strict} = 0$), weakened ($\zeta_{Strict,Mild} = 0$), or eliminated altogether ($\zeta_{Mild,Pre} = \zeta_{Strict,Pre} = 0$). Thus, only $\zeta_{Pre,Mild}$ and $\zeta_{Pre,Strict}$ need to be estimated, as they can be used to directly calculate the expectation of no state change $\zeta_{Pre,Pre} = 1 - \zeta_{Pre,Mild} - \zeta_{Pre,Strict}$. Our procedure can be easily extended to allow for additional transitions between the states.[16]

We estimate the full model using SMM to obtain a value of each parameter in ψ. We simulate an economy of 457 firms (matching the number in our sample), whose optimal choices are computed at a daily frequency for 14 years. We choose an initial burn-in period of five years, during which the economy remains in the *Pre* state (regulatory state changes can occur starting in year 6). We simulate this economy $S = 100$ times in each set of simulations.

[16]For example, it is reasonable to assume that an agreement to impose moderate emissions penalties may later be abandoned. This would require adding only one parameter $\zeta_{Mild,Pre}$ to ψ.

To match the frequency of the empirical data, we cumulate the daily outcomes from each simulation to a quarterly time period when constructing the simulated moment vector. In the SMM objective function Eq. (IA.3) that updates the values in $\hat{\psi}$, the weighting matrix W is the inverse of the covariance matrix of the empirical moments, estimating using the influence function approach (Erickson and Whited, 2000). Table 4 lists all of the moments used to identify parameters in ψ.

Following Section 3.2, we map the Paris Agreement to the *Strict* state during estimation, and designate the *Mild* state as a counterfactual that did not occur. The reason is that the signatories agreed to a goal of limiting temperature increases to 2℃, which requires emissions penalties in line with the higher end of the range of SCC estimates in Nordhaus (2019). Moreover, few analysts or market participants indicated that they were anticipating a significantly more ambitious set of climate goals than COP 21 produced. We discard simulations in which the *Mild* state arises or no state change occurs. However, simulated model outcomes in the *Pre* state still depend in part on $\zeta_{Pre,Mild}$ and the financing costs $\phi_{G,Mild}$ and $\phi_{B,Mild}$, since firms optimize in that period while taking account that the *Mild* state may arise.

4.2.2 Pre-set parameters

We pre-set other model parameters by choosing values that are consistent with BCW and most of the empirical literature. Specifically, we set the risk-free rate to $r_f = 5\%$, the carry cost of cash to $\lambda = 1.5\%$, the capital depreciation rate to $\delta = 15\%$, and the centering parameter of the investment adjustment cost function to $\nu = 15\%$. The market price of productivity risk is set to $\eta = 0.4$, based on the average Sharpe ratio for U.S. equities from 1961 to 2017. Following BCW, the correlation between firm and aggregate cashflow shocks equals $\rho = 0.4$. We set the liquidation value to $\omega = 1$.

We calibrate the carbon tax rates to $\tau_{B,Mild} = 0.19\%$ and $\tau_{B,Strict} = 2\%$, based on SCC estimates in Nordhaus (2019). Specifically, we first compute the average Scope 1 emissions of the treated firms in our sample per dollar of assets. To obtain $\tau_{B,Mild}$ we multiply this amount by \$75.5, which is the midpoint of the SCC's range of \$43–\$108 per ton for less ambitious climate targets. Similarly, $\tau_{B,Strict}$ equals average emissions multiplied by \$218.5, the midpoint of the SCC's range of \$158–\$279 per ton for limiting temperature increases to 2℃.

We set the κ_s parameters to $\kappa_{Pre} = log(3)$, $\kappa_{Mild} = -log(3)$ and $\kappa_{Strict} = -log(3)$. These values are based on BCW, and they represent that investors are averse to the risk of a state change that introduces more regulation and higher equity issuance costs for brown firms. The average cashflow growth rate is set to $\mu = 0.26$ and cashflow volatility is set to $\sigma = 0.28$, based on estimates from Gryglewicz et al. (2022). To obtain these values, we match firms across our and their samples (with a successful match for 348 out of 457 firms), and then compute the average cashflow growth rate and cashflow volatility for the matched firms.

4.3 Identification of unknown model parameters

4.3.1 Transition intensities

We integrate our structural estimation with a DiD model that closely follows BK:

$$Y_{i,t} = \beta_0 + \beta_1 \times Treated \times Post + \beta_2 \times Treated + \beta_3 \times Post + \gamma X_{i,t} + \psi_i + \epsilon_{i,t}, \quad (18)$$

We estimate Eq. (18) using simulated and actual data for three quarterly outcome variables $Y_{i,t}$: Stock returns, investment, and cash holdings. $Treated$ equals 1 for brown (high-emissions) firms and 0 for green (low-emissions) firms. $Treated2$ is defined similarly, but brown firms in the empirical data are classified based on the 75^{th} percentile of emissions in 2014. When the model is estimated with simulated data, $Post$ equals 1 for the four quarters after a state change to $Strict$ in the simulation, and 0 for the four quarters preceding the state change (during the initial Pre period). When it is estimated using actual data, $Post$ equals 1 for Q1 through Q4 of 2016, and 0 for Q4 of 2014 through Q3 of 2015. All regressions also include lagged values of the investment and cash variables as controls in their respective regressions, given their high degree of persistence.[17]

Table 2 shows regression results estimated using actual data. The interaction-term coefficients are highly statistically significant in all specifications. In columns (1) and (2), the positive coefficient indicates that after the Paris Agreement, stock returns of brown firms rose relative to the returns of green firms. This is consistent with the evidence in BK. In columns (3) and (4), the negative coefficient implies that brown firms significantly cut invest-

[17]Controls for the investment regression include lags of investment up to three quarters before. The cash regression includes the one quarter lag of the cash ratio as a control.

ment relative to green firms after the agreement. Finally, the coefficient on cash in columns (5) and (6) is positive, consistent with tighter financing conditions and an increased need for precautionary savings for brown relative to green firms after Paris.

BK interpret the relative increase in brown firms' stocks after the Paris Agreement as evidence that investors became more aware about these firms' exposures to the risk of future climate-associated regulation. In contrast, the market price of risk in Eq. (5) of our model is not state-contingent. Instead, brown firms' required returns increase after a change to the *Strict* state because their expected future profitability decreases while financing costs rise. The increase in returns also reflects the positive effect of a resolution of uncertainty regarding the climate regulation state, to which brown firms are more sensitive.

To verify our assumption that a state change from *Pre* to *Mild* is a counterfactual outcome, we re-estimate Eq. (18) using simulated data with *Post* set to 1 when the *Mild* state arises, and compare the estimates to simulated regressions with *Post* equal to 1 for a change to the *Strict* state. We do this separately with $\zeta_{Pre,Mild}$ and $\zeta_{Pre,Strict}$ each calibrated to 0.3, and with both values set to 0 (i.e., an absence of anticipation). The results in Table IA.2 show that brown firms increase investment after a change to the *Mild* state that is partially anticipated, but cut investment after a change to the *Strict* state regardless of anticipation. Thus, the investment cut documented in Table 2 is more consistent with firms viewing the Paris Agreement as a strict rather than mild outcome.

Separately, we verify that the DiD interaction term coefficients are indeed informative of the unknown transition intensity parameters by conducting a set of comparative statics exercises. These exercises take the final parameter values obtained from the structural estimation in Section 4.4, re-solve the model repeatedly by changing only $\zeta_{Pre,Mild}$ or $\zeta_{Pre,Strict}$, and re-estimate Eq. (18) using the resulting simulated dataset.

Figure 1 presents the estimates for the interaction coefficient β_1, with separate plots for each outcome variable and transition intensity. Overall, the simulated DiD coefficient varies monotonically with most values of the transition intensity parameters. Importantly for identification, increases in $\zeta_{Pre,Mild}$ and $\zeta_{Pre,Strict}$ have opposite effects on the coefficient estimates, which is consistent with the intuition of how anticipation affects firm outcomes. Specifically, when $\zeta_{Pre,Mild}$ is high, brown firms anticipate that the *Pre* state is substantially more likely to change to $Mild$ than to $Strict$, and that the future carbon tax is likely to be small. A realization of the $Strict$ state comes as more of a surprise when $\zeta_{Pre,Mild}$ is high than when it is low, and brown firms respond by reducing investment by a larger amount relative to green firms. The opposite is true when $\zeta_{Pre,Strict}$ is high: Brown firms expect a future carbon tax to be high, and already begin to adjust investment levels during the Pre state. Upon a realization of the $Strict$ state, the additional reduction in investment is smaller when $\zeta_{Pre,Strict}$ is high than when it is low. For similar reasons, brown firms increase cash holdings by a larger amount when the realization of the $Strict$ state comes as more of a surprise (i.e., when $\zeta_{Pre,Mild}$ is higher or $\zeta_{Pre,Strict}$ is lower), and their valuations also rise more.

4.3.2 Investment adjustment costs

The primary moments used to identify θ_G and θ_B are the average investment levels for green and brown firms over the whole sample period. This follows from Eq. (17), which shows that an increase in investment adjustment costs leads to a decrease in optimal investment. The interaction term coefficient from Eq. (18), using investment as the outcome variable, also provides additional identifying information. For example, when θ_B is high, brown firms' investment levels should decrease less than those of green firms after a carbon tax is introduced. Because a DiD model only examines relative changes in firm outcomes, the regression estimates cannot be used to separately identify θ_G and θ_B. Further identifying information comes from moments constructed using market-to-book ratios.

4.3.3 Financing costs

We use eight moments that measure the frequency and amount of equity issuance to identify the unknown financing cost parameters. The fixed financing costs ($\phi_{G,Pre}$, $\phi_{G,Mild}$, $\phi_{G,Strict}$, $\phi_{B,Pre}$, $\phi_{B,Mild}$, and $\phi_{B,Strict}$) should affect the frequency of equity issuance. The model allows these costs to vary across states and firm types, so we construct four separate moments based on the issuance frequency of green and brown firms, before and after the adoption of a carbon tax. The marginal financing costs (γ_G and γ_B) should affect the amount of equity that firms issue. These costs do not vary by state in our model. However, equity issuance amounts depend on both marginal and fixed financing costs, and the latter are state-contingent in the model. Thus, we also construct four separate moments based on the average amount of issuance by green and brown firms, before and after the Paris Agreement.

Additional identifying information for each of the financing cost parameters comes from the interaction term coefficient in Eq. (18), using cash holdings as the outcome variable. This is due to the tight link predicted by the model between external financing frictions and precautionary cash savings.

4.4 Structural estimation results

Table 3 presents estimates of the unknown parameters in ψ. The estimated transition probability $\zeta_{Pre,Mild}$ from the Pre to $Mild$ state is 0.51, while the estimated transition probability $\zeta_{Pre,Strict}$ to the $Strict$ state is 0.26. This implies that the estimated likelihood $\zeta_{Pre,Pre}$ of no state change is 0.23. Taken together, firm outcomes observed in the actual data are consistent with market participants anticipating that some type of emissions penalties would be agreed upon at the COP 21 meeting, but also underestimating the extent of the penalties. In particular, observed outcomes indicate that prior to the COP 21 meeting, market participants anticipated that mild emissions penalties were twice as likely to be adopted as strict.

Other estimates from Table 3 indicate that the investment adjustment cost of green firms is slightly smaller than that of brown firms (θ_G of 33.615, compared to θ_B of 40.138). This is consistent with average investment levels in the actual data, which are similar between the two types of firms. In contrast, brown firms' equity issuance cost parameters are higher in all states than those of green firms. Moreover, the fixed cost of equity issuance for brown firms increases significantly with the degree of potential emissions penalties ($\phi_{B,Pre}$ of 0.5, compared to $\phi_{B,Strict}$ of 0.069), implying that financing terms have tightened for these firms since the Paris Agreement. The opposite is true for green firms, although differences in the fixed cost of financing are small across states ($\phi_{G,Pre}$ of 0.049, compared to $\phi_{G,Strict}$ of 0.047).

Table 4 reports the 17 moments that SMM attempts to match during model estimation, and shows both the simulated and empirical moment values. The DiD coefficient estimates obtained using model-simulated data closely match the corresponding estimates obtained using actual data, for all three outcome variables. This is consistent with the small standard errors for the estimated transition intensity parameters in Table 3. The model also does a reasonable job of matching the empirical equity issuance cost moments. Simulated average investment levels are higher than the empirical moments. This could be because the model does not account for investment irreversibility, which would imply a higher hurdle rate before a firm undertakes any investment. Finally, the simulated moments related to the market-to-book ratio are close to the empirical moments in order of magnitude, although the simulated moments decrease after the Paris Agreement while the empirical moments increase.

4.5 Counterfactual analysis

To gauge the magnitude of how anticipation affected firm outcomes, we solve firms' optimal responses to a change in the regulatory policy state in the counterfactual absence of any anticipation. Specifically, we run a set of model simulations with the transition intensities $\zeta_{Pre,Mild}$ and $\zeta_{Pre,Strict}$ both set to 0 (so firms assign a probability of 1 to staying in the *Pre* state), and with other unknown model parameters set to the estimated values from Table 3. In this counterfactual, solutions for each possible policy state are effectively independent of each other. The model begins in the *Pre* state in the initial time period, and we then artificially impose a change to the *Strict* state in a future time period (necessary since the counterfactual Markov chain matrix assigns a 0 probability to such a transition). We then re-estimate Eq. (18) using the counterfactual panel dataset obtained from this set of simulations. This exercise is analogous to the reduced-form analysis of a natural experiment in which the exogenous policy change is completely unanticipated, and it provides a quantitative estimate of the causal effect of the Paris Agreement on firm outcomes.

Table 5 compares the DiD coefficients obtained from estimating Eq. (18) on the counterfactual dataset to the DiD coefficients reported in Table 3 (estimated on the dataset from model simulations with non-zero transition intensities). The counterfactual DiD coefficient on investment is -0.0022, compared to the simulated DiD coefficient of -0.0015. This indicates that in the presence of anticipation, the average decline in the investment of brown relative to green firms is about 35% smaller than the true causal effect. Because brown firms anticipate some form of future emissions penalties, they invest less than they optimally would (relative to green firms) in the absence of anticipation. Therefore, the realization of the *Strict* state results in smaller investment changes than would arise in the absence of anticipation.

The counterfactual DiD coefficient on cash of 0.0158 is 50% larger than the simulated coefficient of 0.0079. This indicates that the causal effect of the Paris Agreement on cash holdings is bigger than firms' responses in the presence of anticipation. Because firms anticipate a future policy state with higher external financing costs, they already engage in precautionary cash savings in the Pre state. The effect of anticipation on cash is larger than the effect on investment, perhaps because firms incur relatively lower costs to adjust cash holdings than investment. Finally, the DiD coefficient on returns of 0.0814 is 3% larger than the coefficient measured in the counterfactual without anticipation.

5 Conclusion

Many empirical methods struggle to accurately estimate the causal effects of policy events when agents adjust their behavior in anticipation of future policy changes—a phenomenon known as anticipation bias. We propose an approach that improves causal inference by explicitly incorporating agents' beliefs into estimation, combining reduced-form and structural techniques around the observed outcomes of a single policy shift. We demonstrate the value and application of this method using the Paris Agreement, a widely studied event linked to increased climate regulatory risk. Our analysis reveals that anticipation can distort not only the magnitude but also the direction of estimated treatment effects on firm risk when using standard models such as difference-in-differences. We provide clear, practical guidance for addressing the gap between true causal effects and conventional estimates.

References

Albuquerque, R., M. S. Kruttli, Y. Ma, and S. W. Watugala (2025). Financing and the green paradox. *SSRN Working Paper*.

Alessi, L., S. Battiston, and V. Kvedaras (2024). Over with carbon? Investors' reaction to the Paris Agreement and the US withdrawal. *Journal of Financial Stability 71*, 101232.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics 151*, 103745.

Baker, A. C., D. F. Larcker, and C. C. Wang (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics 144*, 370–395.

Bolton, P., H. Chen, and N. Wang (2011). A unified theory of tobin's q, corporate investment, financing, and risk management. *Journal of Finance 66*(5), 1545–1578.

Bolton, P., H. Chen, and N. Wang (2013). Market timing, investment and risk management. *Journal of Financial Economics 109*, 40–62.

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. *Journal of Finance 78*(6), 3051–3757.

Borochin, P. and J. Golec (2016). Using options to measure the full value-effect of an event: Application to obamacare. *Journal of Financial Economics 120*, 169–193.

Briggs, J., D. Cesarini, E. Lindqvist, and R. Östling (2021). Windfall gains and stock market participation. *Journal of Financial Economics 139*, 57–83.

Bustamante, M. C. and F. Zucchi (2024). Dynamic carbon emission management. ECB Working Paper No. 2024/2885.

Cronqvist, H., T. Ladika, E. Pazaj, and Z. Sautner (2024). Limited attention to detail in financial markets: Evidence from reduced-form and structural estimation. *Journal of Financial Economics 154*, 103811.

Degryse, H., R. Goncharenko, C. Theunisz, and T. Vadasz (2023). When green meets green. *Journal of Corporate Finance 78*, 102355.

Döttling, R. and M. Rola-Janicka (2025). Too levered for Pigou: Carbon pricing, financial constraints, and leverage regulation. *Journal of Financial Economics*, forthcoming.

Eisfeldt, A. and G. Schubert (2024). AI and finance. NBER Working Paper 33076.

Eisfeldt, A., G. Schubert, M. B. Zhang, and B. Taska (2025). Generative AI and firm values. NBER Working Paper 31222.

Erickson, T. and T. M. Whited (2000). Measurement error and the relationship between investment and Q. *Journal of Political Economy 108*, 1027–1057.

Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of Financial Economics 116*(1), 1–22.

Garel, A., A. Romec, Z. Sautner, and A. Wagner (2024). Do investors care about biodiversity? *Review of Finance 28*(4), 1151–1186.

Ginglinger, E. and Q. Moreau (2023). Climate risk and capital structure. *Management Science 69*(12), 7492–7516.

Gryglewicz, S., L. Mancini, E. Morellec, E. Schroth, and P. Valta (2022). Understanding cash flow risk. *Review of Financial Studies 35*, 3922–3972.

Hayashi, F. (1982). Tobin's marginal q and average q: a neoclassical interpretation. *Econometrica 50*, 215–224.

Heath, D., M. C. Ringgenberg, M. Samadi, and I. M. Werner (2023). Reusing natural experiments. *Journal of Finance 78*, 2329–2364.

Hennessy, C. and D. Livdan (2021). Learning, parameter drift, and the credibility revolution. *Journal of Monetary Economics 117*, 395–417.

Hennessy, C. and I. Strebulaev (2020). Beyond random assignment: Credible inference and extrapolation in dynamic economies. *Journal of Finance 75*(2), 825–866.

Hou, K., C. Xue, and L. Zhang (2015). Digesting anomalies: An investment approach. *Review of Financial Studies 28*(3), 650–705.

Klausmann, J., P. Krueger, and P. Matos (2024). The green transition: Evidence from corporate green revenues. Working Paper.

Nordhaus, W. (2019). Climate change: The ultimate challenge for economics. *American Economic Review 109*(6), 1991–2014.

Pang, J., Z. Liu, W. H. Hou, and Y. Tao (2023). How does the paris agreement affect firm productivity? international evidence. *Finance Research Letters 56*(104150).

Pedersen, L. H. (2024). Carbon pricing versus green finance. *Working Paper*.

Pástor, and P. Veronesi (2013). Political uncertainty and risk premia. *Journal of Financial Economics 110*, 520–545.

Seltzer, L., L. T. Starks, and Q. Zhu (2025). Climate regulatory risks and corporate bonds. FRB of New York Staff Report No. 1014.

Subramanian, A. (2004). Option pricing on stocks in mergers and acquisitions. *Journal of Finance 59*(2), 795–829.

Whited, T. (2023). Integrating structural and reduced-form methods in empirical finance. *Journal of Financial Econometrics 21*, 597–615.

Worland, J. (2015). Kerry urges strong agreement on climate change ahead of Paris conference. *TIME*. October 18, 2025.

Zhang, S. (2025). Carbon returns accross the globe. *Journal of Finance 80*(1), 615–645.

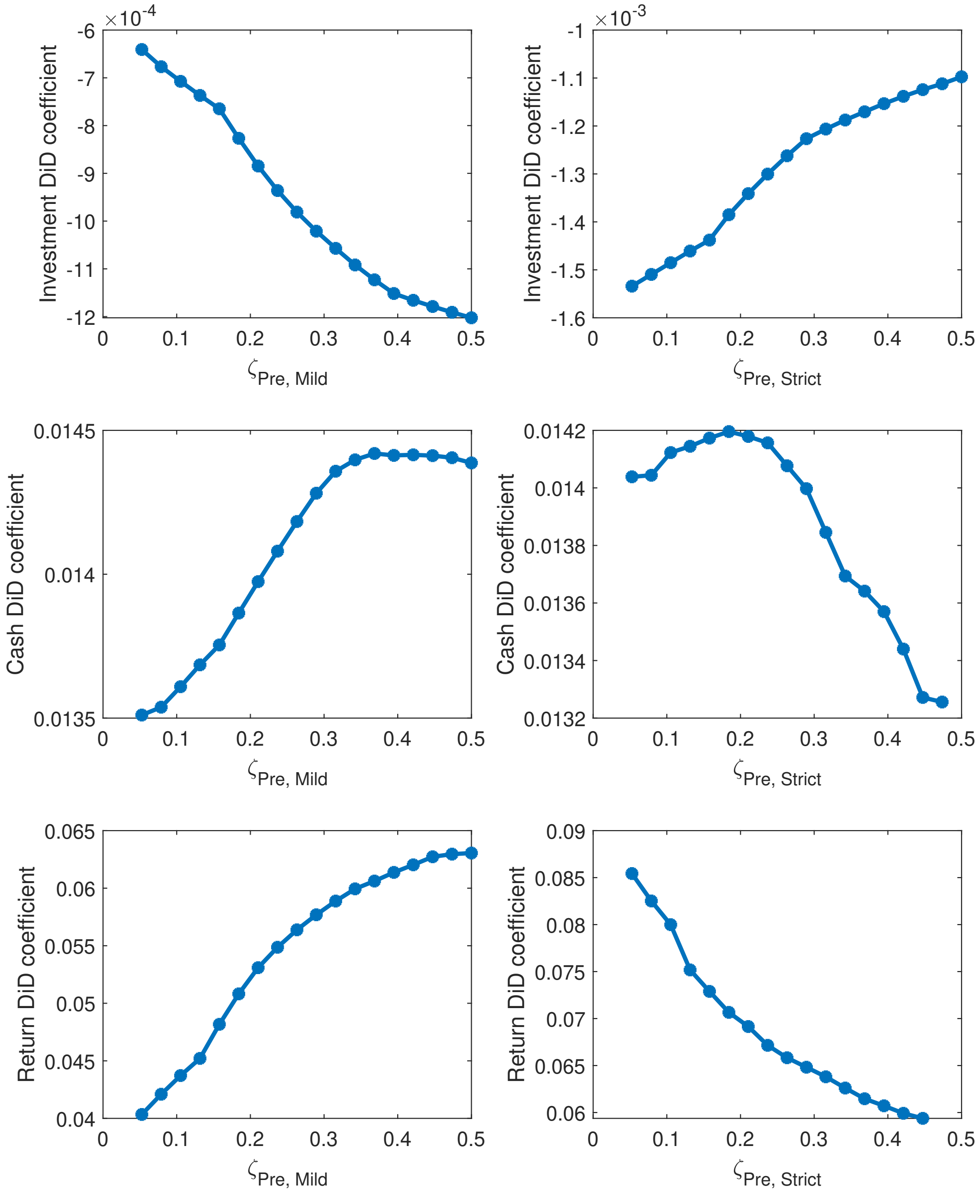

Figure 1: Comparative statics. This figure reports coefficient estimates for $Treated \times Post$, obtained by estimating Eq. (18) using a dataset of simulated model outcomes. For each panel, a set of 100 simulations is run on the full model from Section 2.3, with the transition intensity $\zeta_{Pre,Mild}$ or $\zeta_{Pre,Strict}$ set to the corresponding value on the x-axis. All other model parameters are calibrated to reasonable values. In the figures on the left, *Post* in Eq. eqrefeq:DiD˙model equals 1 for time periods in which the model is in the *Mild* state is realized, and 0 for time periods in which the model is in the *Pre* state. In the figures on the right, *Post* is similarly defined but equals 1 for time periods in which the model is in the *Strict* state.

Table 1: Summary statistics. This table provides summary statistics calculated on actual data, for the key variables used in the empirical analysis. The sample is 457 U.S. firms in the intersection of Compustat, CRSP, and S&P Trucost, and covers the period from Q1 of 2015 through Q4 of 2016. *Investment* represents quarterly investment computed using quarterly differences in capital expenditures (Compustat data item $CAPXY$) scaled by total assets (ATQ); *Cash* is the ratio of cash and cash equivalents ($CHEQ$) to total assets (ATQ); *Returns* represent raw stock returns cumulated from monthly to quarterly; *Equity issuance* denotes the amount of equity issued over each calendar quarter, computed as $STTKQ$ scaled by total assets (ATQ). The statistics on *Equity issuance* are conditional on issuance. *Market-to-Book* is computed as the ratio of the stock market capitalization of the firm (price per share at the end of the quarter times number of shares outstanding ($PRCCQ \times CSHOQ$) plus the book values of long term and short term debt ($DLTTQ$ and $DLCQ$), to the sum of the book values of long-term and short-term debt and book value of equity ($CEQQ$). *Mom* is the cumulative stock return over the previous year; *Beta* is the average market beta over the calendar quarter; *Volatility* is average monthly stock return volatility over the quarter; *Sales growth* is the year-on-year change in quarterly sales $SALEQ$ over market capitalization; *EPS growth* is the year-on-year change in diluted EPS excluding extraordinary items ($EPSFXQ$) over the price per share ($PRCQQ$); *Size* is the log of the total equity market capitalization; *PPE* is the logarithm of property, plant and equipment ($PPENTQ$). All variables, with the exception on *Equity issuance*, are winsorized at the top and bottom 2.5%. *Equity issuance* is truncated at values above 50% to avoid the inclusion of merger and acquisition related stock issues. *Equity issuance* is truncated at the bottom 1% to avoid stock issues related to the exercise of employee stock options.

Variable	Mean	SD	25^{th}	50^{th}	75^{th}	N
Investment	0.016	0.013	0.006	0.012	0.020	3,766
Cash	0.088	0.085	0.020	0.063	0.129	3,966
Returns	0.007	0.186	-0.08	0.007	0.102	3,998
Equity issuance	0.020	0.029	0.004	0.006	0.017	3,467
Market-to-Book	1.898	1.115	1.140	1.563	2.282	3,860
Mom	0.018	0.088	-0.02	0.019	0.064	3,958
HHI	0.250	0.149	0.142	0.227	0.333	3,838
Beta	0.254	0.147	0.139	0.233	0.354	3,959
Volatility	0.093	0.058	0.056	0.074	0.108	3,962
Sales growth	-0.02	0.081	-0.03	-0.00	0.008	3,960
EPS growth	-0.00	0.049	-0.00	0.000	0.004	3,974
Size	9.131	1.613	8.114	9.045	10.21	3,973
PPE	8.447	1.425	7.443	8.408	9.379	3,943

Table 2: DiD regressions in actual data. This table reports results from DiD regression Eq. (18) estimated using actual data. The sample is 457 U.S. firms in the intersection of Compustat, CRSP, and S&P Trucost, and covers the period from Q1 of 2015 through Q4 of 2016. The dependent variable is *Stock Returns* in columns (1) and (2), *Investment* in columns (3) and (4), and *Cash* in columns (5) and (6). *Post* equals 1 for the Q1 through Q4 of 2016, and 0 for Q1 through Q3 of 2015 (Q4 of 2015 is omitted). *Treated* equals 1 for firms whose ratio of Scope 1 emissions to revenues in 2014 ranks above the 66$^{\text{th}}$ percentile of the sample distribution, and 0 for a matched set of firms whose emission intensities are below the 66$^{\text{th}}$ percentile. *Treated2* is similarly defined, except based on the 75$^{\text{th}}$ percentile of emission intensities in 2014. Control variables include lagged values of $Size$, $Market - to - Book$, $Leverage$, $Cash$, ROE, PPE, *and* $Investment$. *All specifications include firm fixed effects. Standard errors are presented in parentheses. ***, **, * indicate significance levels of 1%, 5%, and 10%, respectively.*

Outcome	*Stock Returns*		*Investment*		*Cash*	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × *Post*	0.0813***		-0.0014**		0.0066**	
	(0.0124)		(0.0005)		(0.0020)	
Treated2 × *Post*		0.0785***		-0.0015***		0.0057**
		(0.0131)		(0.0005)		(0.0020)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	3,331	3,155	3,111	2,947	3,268	3,092
R-squared	0.0729	0.0763	0.1412	0.1447	0.1720	0.1627

Table 3: Parameter estimates. This table presents estimates for each unknown parameter in the set ϕ, obtained by estimating the full model from Section 2.3 following our integration procedure described in Section 4. The model is estimated using SMM. Standard errors are presented in parentheses.

Description	Symbol	Estimate
(1) Transition intensity from *Pre* state to *Mild* state	$\zeta_{Pre,Mild}$	0.5123
		(0.0599)
(2) Transition intensity from *Pre* state to *Strict* state	$\zeta_{Pre,Strict}$	0.2602
		(0.0029)
(3) Investment adjustment costs for green firms	θ_G	33.6150
		(2.7066)
(4) Investment adjustment costs for brown firms	θ_B	40.1377
		(9.3956)
(5) Marginal cost of equity issuance for green firms	γ_G	0.0103
		(0.0094)
(6) Fixed cost of equity issuance for green firms in *Pre* state	$\phi_{G,Pre}$	0.0490
		(0.1050)
(7) Fixed cost of equity issuance for green firms in *Mild* state	$\phi_{G,Mild}$	0.0480
		(0.0956)
(8) Fixed cost of equity issuance for green firms in *Strict* state	$\phi_{G,Strict}$	0.0470
		(0.4119)
(9) Marginal cost of equity issuance for brown firms	γ_B	0.1351
		(0.0086)
(10) Fixed cost of equity issuance for brown firms in *Pre* state	$\phi_{B,Pre}$	0.0500
		(0.1837)
(11) Fixed cost of equity issuance for brown firms in *Mild* state	$\phi_{B,Mild}$	0.0585
		(0.1605)
(12) Fixed cost of equity issuance for brown firms in *Strict* state	$\phi_{B,Strict}$	0.0690
		(0.5001)

Table 4: Simulated and empirical moments. This table presents the simulated and empirical that we use to estimate the full model from Section 2.3 following our integration procedure described in Section 4. The model-simulated moments are calculated using the unknown parameter estimates reported in Table 3. Empirical moments are calculated using a sample of 457 U.S. firms in the intersection of Compustat, CRSP, and S&P Trucost, covering the period from Q1 of 2015 through Q4 of 2016. Standard errors are presented in parentheses.

Description	Simulated	Empirical
Panel A. Difference-in-difference moments		
(1) $\beta_{1,Ret}$	0.0814	0.0813
	(0.000)	(0.000)
(2) $\beta_{1,Inv}$	-0.0015	-0.0014
	(0.000)	(0.0000)
(3) $\beta_{1,Cash}$	0.0079	0.0066
	(0.000)	(0.0000)
Panel B. Equity issuance moments		
(4) Issuance frequency green firms before	0.0673	0.0585
	(0.0027)	(0.0023)
(5) Issuance frequency brown firms before	0.0766	0.0637
	(0.0029)	(0.0026)
(6) Issuance frequency green firms after	0.0879	0.0539
	(0.0030)	(0.0019)
(7) Issuance frequency brown firms after	0.1034	0.0722
	(0.0030)	(0.0029)
(8) Issuance amount green firms before	0.0100	0.0137
	(0.0002)	(0.0002)
(9) Issuance amount brown firms before	0.0100	0.0238
	(0.0002)	(0.0004)
(10) Issuance amount green firms after	0.0300	0.0129
	(0.0005)	(0.0002)
(11) Issuance amount brown firms after	0.0400	0.0267
	(0.0006)	(0.0004)
Panel C. Investment moments		
(12) Investment green firms	0.0392	0.0120
	(0.000)	(0.0003)
(13) Investment brown firms	0.0358	0.0188
	(0.000)	(0.0004)
Panel D. M/B moments		
(14) M/B green firms before	1.5233	1.4381
	(0.0017)	(0.0183)
(15) M/B green firms after	1.4846	1.7732
	(0.0017)	(0.0183)
(16) M/B brown firms before	1.1522	1.3095
	(0.0017)	(0.0130)
(17) M/B brown firms after	0.886	1.7087
	(0.0019)	(0.0130)

Table 5: DiD moments in simulated counterfactual data. This table compares coefficient estimates for $Treated \times Post$ obtained by estimating Eq. (18) using two datasets of simulated model outcomes. In the first row the dependent variable is $Investment$, in the second row it is $Cash$, and in the third row it is *Stock Returns*. In the "Counterfactual" column, the dataset is constructed by running a set of 100 simulations of the full model from Section 2.3, with $\zeta_{Pre,Mild}$ and $\zeta_{Pre,Strict}$ are set to 0 (to simulate a counterfactual economy with no anticipation of a regulatory state change), while all other unknown parameter values are set to the values reported in Table 3. In the "Estimated" column the simulated dataset is constructed in the same way, except $\zeta_{Pre,Mild}$ and $\zeta_{Pre,Strict}$ are also set to the values reported in Table 3. The column Δ (%) reports the percentage difference between the Counterfactual and Matched coefficient estimates.

Moment	Counterfactual	Estimated	Δ (%)
$\beta_{1,Inv}$	-0.0022	-0.0015	34.8
$\beta_{1,Cash}$	0.0158	0.0079	50.2
$\beta_{1,Ret}$	0.0787	0.0814	-3.3

Internet Appendix

for

Addressing Anticipation Effects in Finance

IA.1 Calibration

In this internet appendix, we provide a calibration the model with the goal to show how measured treatment effects can deviate from causal effects. Notably, we confirm that the size and sign of the bias depends on which state is realized in the data. Specifically, the size of the bias varies significantly with model parameter estimates, especially the measured transition intensities. This highlights the necessity of obtaining precise parameter estimates when the goal is to quantify causal effects.

We necessarily consider a finite number of possible climate regulation states. To highlight the key implications of anticipation, in this calibration we consider only three states. State 1 features zero carbon taxes and no difference in the costs of capital between green and brown firms. This can be thought of as a *Pre*-Paris Agreement state. We set the fixed costs of equity issuance to $\phi_{B,1} = \phi_{G,1} = 1\%$ and the marginal costs of equity issuance to $\gamma_{B,1} = \gamma_{G,1} = 3\%$ which are generally in line with Bolton et al. (2013). State 2 features a positive carbon tax, set at a relatively low level: $\tau^e_{B,2} = 3\%$.[1] This can be thought of as a possible *Post*-Paris Agreement state. In this state, the cost of equity capital is higher for brown firms compared to green firms. The fixed cost of equity issuance for brown firms is also higher relative to State 1 and increases to $\phi_{B,2} = 1.5\%$. This cost is lower for green firms relative to State 1 and decreases to $\phi_{G,2} = 0.5\%$. The respective marginal costs of equity issuance remain unchanged: $\gamma_{B,2} = \gamma_{G,2} = 3\%$. State 3 represents another possible "*Post*-Paris Agreement" state, whereby compared to State 2, carbon taxes are higher and differences in equity issuance costs between green and brown firms are starker. Namely, the carbon tax is set at a level that is twice as high as the level in State 2: $\tau^e_{B,3} = 6\%$; the fixed cost of equity issuance for green firms is $\phi_{G,3} = 0.1\%$ and the marginal cost is $\gamma_{G,3} = 2\%$; the fixed cost of equity issuance for brown firms is $\phi_{B,3} = 2\%$ and the marginal cost is $\gamma_{B,3} = 3\%$. The technology parameters are set close to the representative firm values in Bolton et al. (2013) and do not vary across climate regulation states.

In the baseline calibration we consider States 2 and 3 as equally likely. We set the transition intensities to these states to: $\zeta_{(1,2)} = \zeta_{(1,3)} = 0.3$. To simplify the analysis, we

[1] Note that the tax for green firms is always zero, hence $\tau^e_{G,2} = 0\%$.

consider these states as absorbing states: once reached, the economy cannot move to any other state (implying $\zeta_{(2,1)} = \zeta_{(2,3)} = \zeta_{(3,2)} = \zeta_{(3,1)} = 0$). The probability of staying in State 1 is $\zeta_{(1,1)} = 1 - \zeta_{(1,2)} - \zeta_{(1,3)} = 0.4$ and the corresponding duration of State 1 is 2.5 years ($1/\zeta_{(1,1)}$). This implies an anticipation of the Paris Agreement of around the two years before it was signed, a time frame that is also largely in line with length of the *Pre* period typically used in DiD studies (such as, for example, in Bolton and Kacperczyk (2023)).

To summarize, all parameters, including financing costs and carbon taxes, are the same for green and brown firms in the *Pre*-Paris Agreement state (State 1). The reasoning for this is to consider brown and green firms in the model as representing matched firms in DiD studies. In the *Post*-Paris Agreement states (States 2 and 3), brown firms face non-zero carbon taxes as well as more expensive external financing, while green firms face cheaper external financing. Table IA.1 summarizes all parameter values.

IA.2 Baseline results

Figure IA.1 shows how average q and investment vary with the cash-to-capital ratio for green and brown firms in each state. Average q is defined similarly to Bolton et al. (2013) as the ratio of enterprise value to capital:

$$q(m, s) = \frac{V(M, K, s) - M}{K} = F(m, s) - m. \tag{IA.1}$$

The first key observation from Figure IA.1 is that even though all model parameters are the same for green and brown firms in State 1, firm values of brown firms are lower compared to green firms. The difference emerges because State 1 valuations incorporate anticipation effects of State 2 and State 3: State 1 firm values of brown firms incorporate the possibility of the state switching to one with positive carbon taxes and higher financing costs, while State 1 values of green firms incorporate the possibility of the state switching to one with lower financing costs.

Panel A shows that average q for brown firms in State 1 is lower compared to average q in State 2. This seems counterintuitive but is directly tied to the fact that State 1 valuations incorporate the possibility of switching to an even worse state than State 2 (State 3). Average q in State 3 is the lowest, in line with this state being absorbing and incorporating the highest level of carbon taxes and highest costs of financing. Panel B shows how investment varies with the cash-to-capital ratio for brown firms in each of the three states. Anticipation effects have a similar impact on the relative levels of investment as on average q.

The effect of anticipation for green firms goes in the opposite direction. Panel C shows that average q in State 1 is higher than average q in State 2, despite cheaper financing costs in State 2. The reason again relates to anticipation: valuations in State 1 reflect the possibility of switching to an even better state relative to State 2 (State 3 where financing costs are lowest). This possibility also affects the levels of investment shown in Panel D, whereby investment in State 1 is higher than investment in State 2. Both valuations and investment in State 3 are higher than in the other states, in line with State 3 incorporating the lowest level of financing costs.

IA.3 Model simulation and DiD analysis

We simulate 100 years of monthly data for 100 green and 100 brown firms. Green and brown firms are subject to both aggregate and idiosyncratic productivity shocks. The climate regulation state represents another aggregate shock that affects all firms. We simulate these states given the transition probability matrix specified in Table IA.1. We impose the simulations to start from State 1, which is considered as the *Pre*-Paris Agreement state. States 2 and 3 represent two possible realizations post the change out of State 1. The realization of each of these states can represent a *Post*-Paris Agreement state.

Figure IA.2 shows how the (cross-sectional) average scaled valuations of green and brown firms change in the months around the change out of State 1. The top panel shows these changes when State 2 is realized; the bottom panel shows the changes in value that occur when State 3 is realized. Consistent with the results in Figure IA.1, green firms are more valuable than brown firms in State 1. Idiosyncratic productivity shocks on average cancel out, and average green and brown firm values, unless there is a change in the climate regulation state, move in tandem due to aggregate productivity shocks. (This is equivalent to satisfying the "parallel trends" assumption in reduced-form empirical work.) When the state switches from 1 to 2, brown firm values increase while green firm values decline. The opposite occurs when State 3 is realized.

We run the following regression model in simulated data:

$$Y_{i,t} = \beta_0 + \beta_1 \times \text{Treated} \times \text{Post} + \beta_2 \times \text{Treated} + \beta_3 \times \text{Post} + \epsilon_{i,t}, \tag{IA.2}$$

where $Y_{i,t}$ denotes the outcome variable in the regression, $Treated$ equals one for brown firms, and zero otherwise, and $Post$ equals one when the state has switched out of State 1, and zero otherwise. We use scaled valuations, investment and returns as outcome variables.

Table IA.2, Panel A, shows the regression results when the model is simulated using the transition intensity matrix in Table IA.1. These results thus incorporate the effects of anticipation on outcome variables. Specifications (1) and (3) show that a change of state from State 1 to State 2 has a significant positive impact on the valuations and investment of brown firms relative to green firms. This is in line with the reasoning in Section IA.2, whereby the realization of State 2 represents an outcome of regulation that is better than expected for brown firms (and worse than expected for green firms), hence the positive coefficient on both valuations and investment. Panel B reports regression results when the model is simulated assuming no anticipation. We do so by using a transition probability matrix whereby all diagonal elements are one and all off-diagonal elements are zero. The signs of the coefficients on $Treated \times Post$ dummies in specifications (1) and (3) are reversed relative to Panel A. Absent anticipation, a change to a state with higher taxes and higher external financing costs leads to lower valuations and lower investment for brown firms relative to green firms.

Specifications (2) and (4) show the effect on valuations and investment of the state switching from State 1 to State 3. The coefficients on $Treated \times Post$ dummies are negative with and without anticipation. However, in the presence of anticipation the effect is smaller compared to the causal effect one obtains from the regressions absent anticipation.

IA.4 Comparative statics

We next conduct a set of comparative statics to examine how treatment and causal effects vary with key model parameters. Figure IA.5 presents comparative statics with respect to carbon tax rates. Panel A shows that when State 2 is realized, the average treatment effect (ATE) declines with the carbon tax rate in State 2 (τ_2^e). The ATE is positive for lower levels of the carbon tax rate, a result that is consistent with the analysis in the previous subsections. However, the ATE turns negative as the carbon tax rate approaches the level of the tax rate in State 3. In other words, the realization of State 2 no longer represents "good news" for brown firms. As the carbon tax τ_2^e increases, the average causal effect (ACE) becomes more negative, consistent with higher taxes reducing firm profitability. The ATE declines at a slower rate compared to the ACE. The reason is that when the carbon tax rate is higher, the anticipation of a potential switch to State 2 has a greater impact on valuations in State 1. When State 2 materializes, the response in valuations is attenuated by previous anticipation. The bias in ATE therefore increases with the carbon tax.

Panel C shows the comparative statics with respect to the carbon tax rate in State 3, τ_3^e, when State 3 is realized. Both average treatment, ATE, and average causal effects, ACE, become more negative as the carbon tax increases. A higher carbon tax in State 3 implies a greater impact of anticipation on valuations in State 1. Similar to before, when State 3 materializes, the response in valuations is attenuated. The bias in ATE again increases with the carbon tax. Panels B and D show the results for Green firms, for which none of the effects change as these firms' capital is assumed not to produce carbon emissions. Figure refFig: comparative statics transition intenstities shows comparative statics with respect to the transition intensities from State 1 to the other two states. Results again indicate that the divergence between ATE and ACE varies significantly with the transition intensities.

These results highlight the importance of precise parameter estimates in determining causal effects and therefore the need for structural estimation. The use of DiD estimates as moments to match in the structural estimation is key for the identification of parameters related to policy regimes. In other words, the integration of reduced-form and structural estimation is necessary to quantify causal effects.

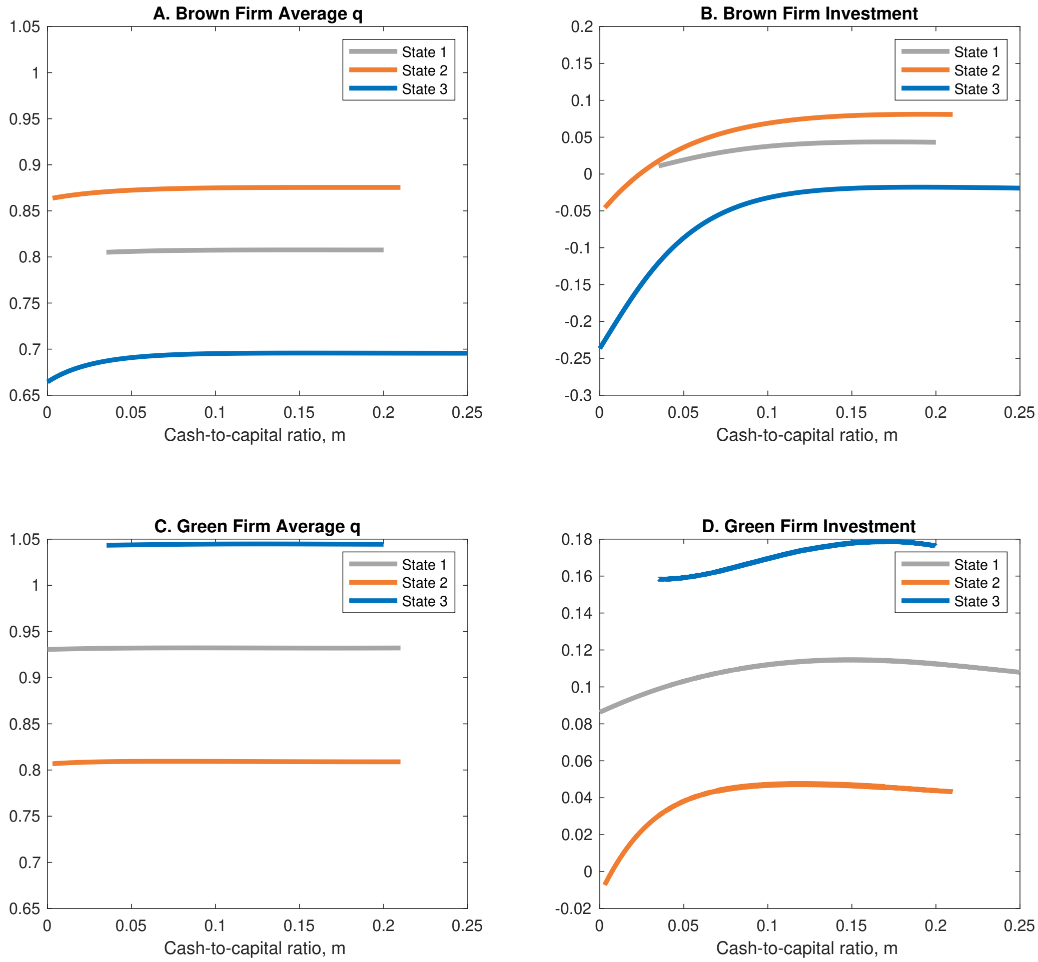

Figure IA.1: Average q and investment. The figures show how average q and investment vary with the cash-to-capital ratio for representative green and brown firms, given the climate regulation state. State 1 denotes the state where there are no carbon taxes. The costs of capital in State 1 are identical for green and brown firms. State 2 denotes a climate regulation state with positive but relatively low carbon taxes. The cost of capital for brown firms in State 2 is higher than in State 1, while the cost of capital for green firms in State 2 is lower than in State 1. State 3 denotes the state with the highest level of carbon taxes. The cost of capital for brown firms in State 3 is higher than in State 2, while the cost of capital for green firms in State 3 is lower than in State 2.

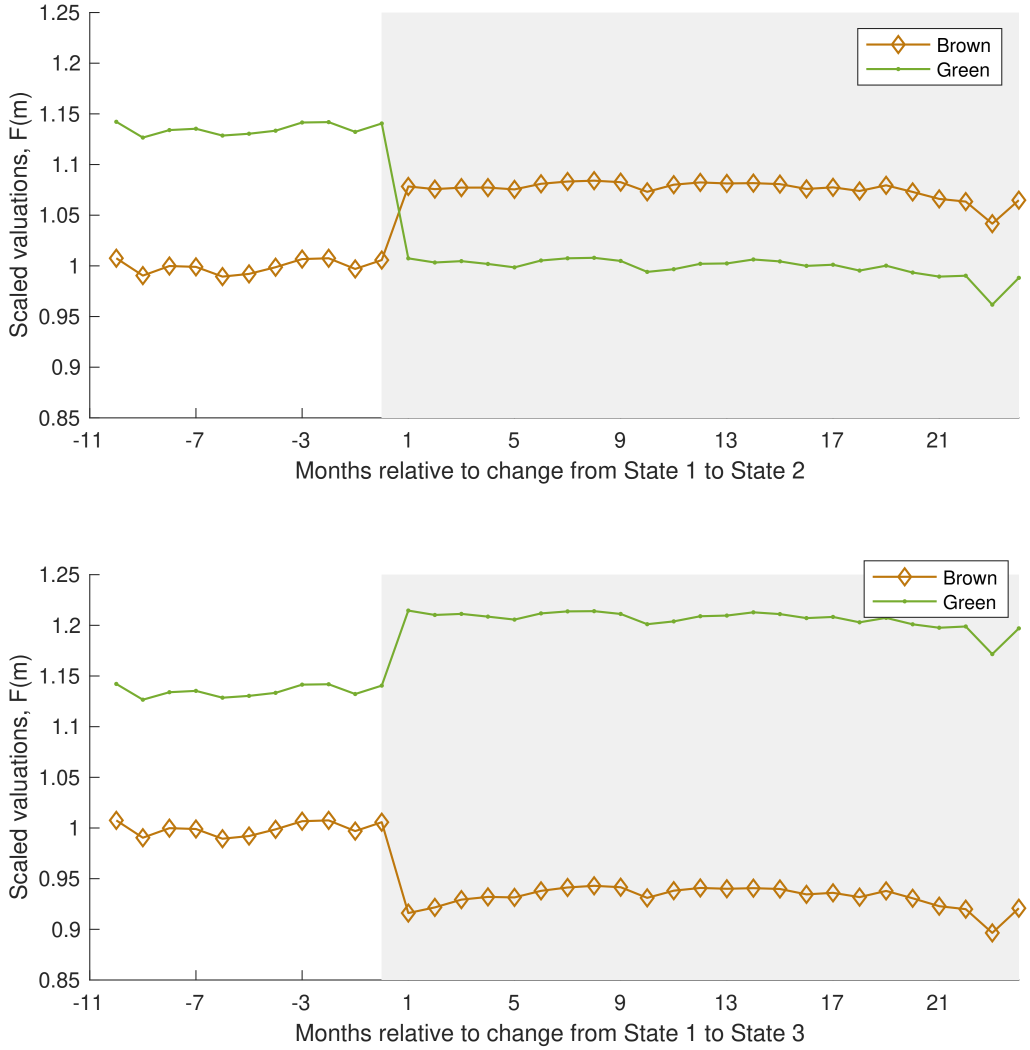

Figure IA.2: Green vs. brown firms. This figure shows how the (cross-sectional) average scaled valuations of green and brown firms change in the months around the change out of State 1. The top panel shows these changes when State 2 is realized; the bottom panel shows the changes in value that occur when State 3 is realized.

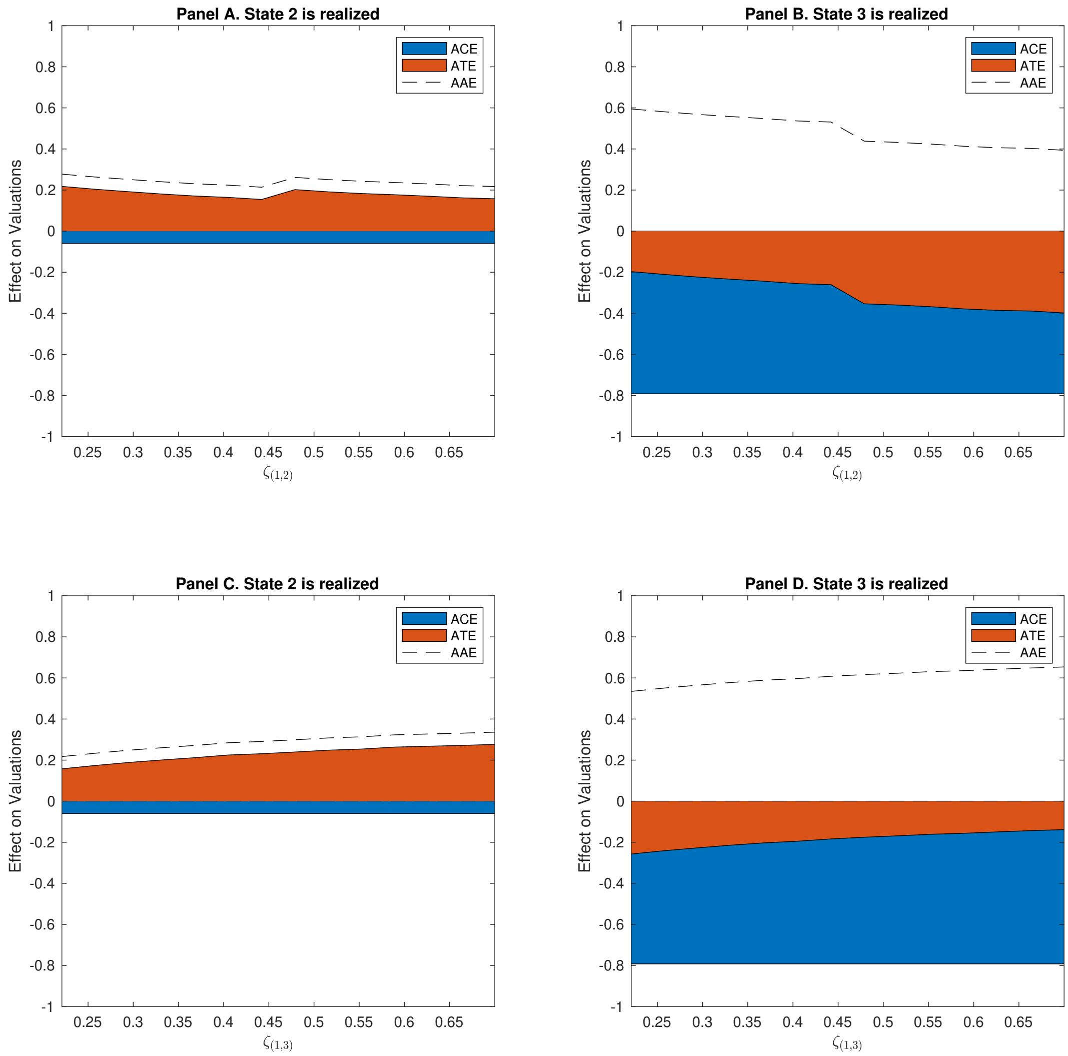

Figure IA.3: Average treatment, causal, and anticipation effects. The figures show how treatment, causal and anticipation effects vary with the transition intensities out of State 1. The top two panels show the variation with respect to the level of the carbon tax in State 2. The first panel shows the effects when State 2 is realized, while the second panel shows the effects when State 3 is realized. The bottom two panels show the corresponding effects when varying the level of the carbon tax in State 3.

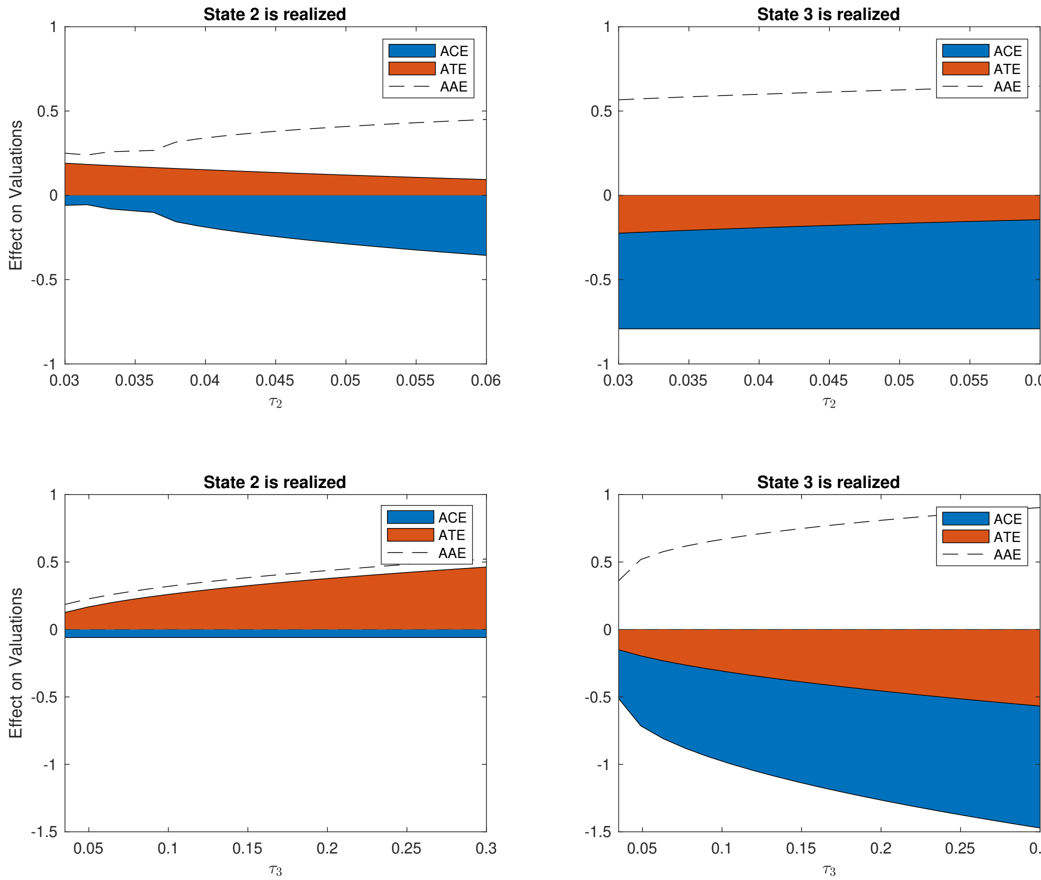

Figure IA.4: Average treatment, causal, and anticipation effects. The figures show how treatment, causal and anticipation effects vary with the level of carbon taxes. The top two panels show the variation with respect to the level of the carbon tax in State 2. The first panel shows the effects when State 2 is realized, while the second panel shows the effects when State 3 is realized. The bottom two panels show the corresponding effects when varying the level of carbon taxes in State 3.

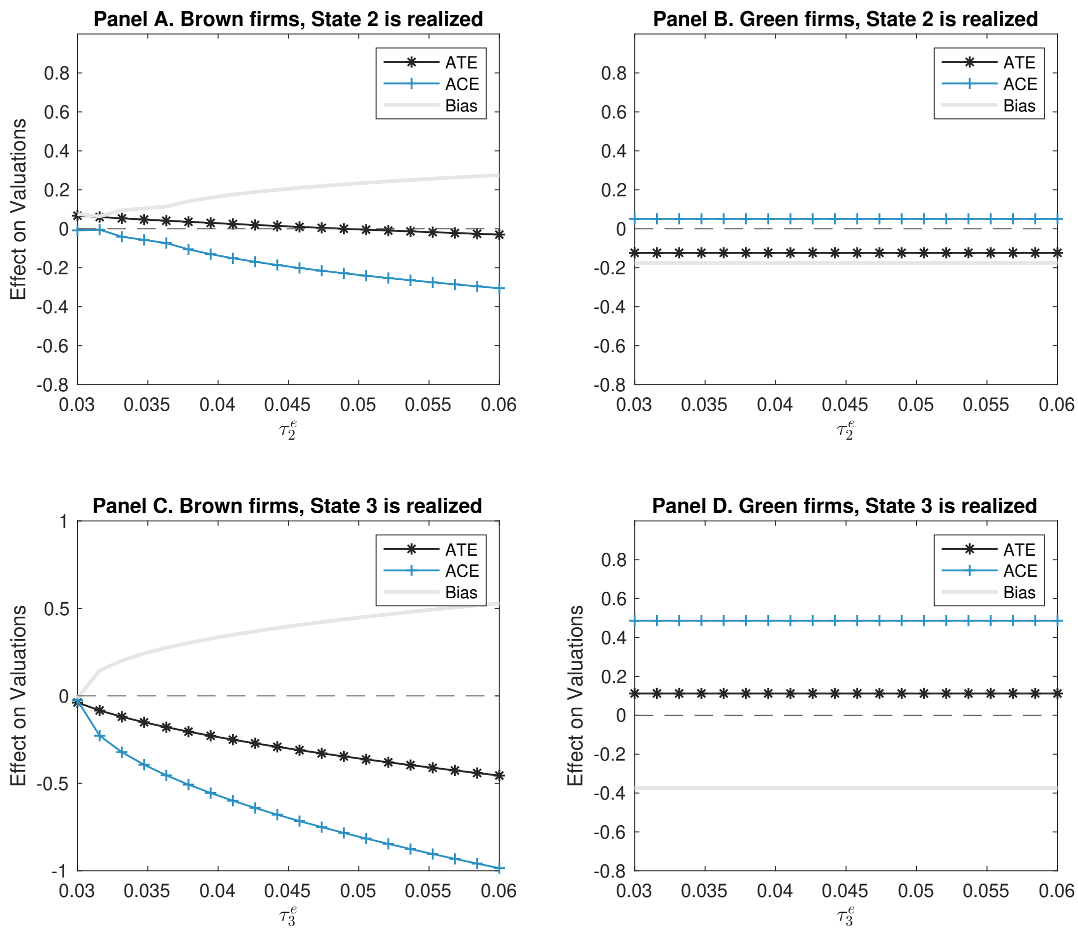

Figure IA.5: Comparative statics: Carbon tax rates. The figures show how the average treatment and causal effects, as well as the bias in treatment effects induced by anticipation vary with the level of carbon taxes. Panels A and B show these comparative statics with respect to the carbon tax rate in State 2, τ_2^e, when State 2 is realized. Panel A shows the effects for brown firms, while panel B shows the effects for green firms. Panels C and D show the comparative statics with respect to the carbon tax rate in State 3, when State 3 is realized.

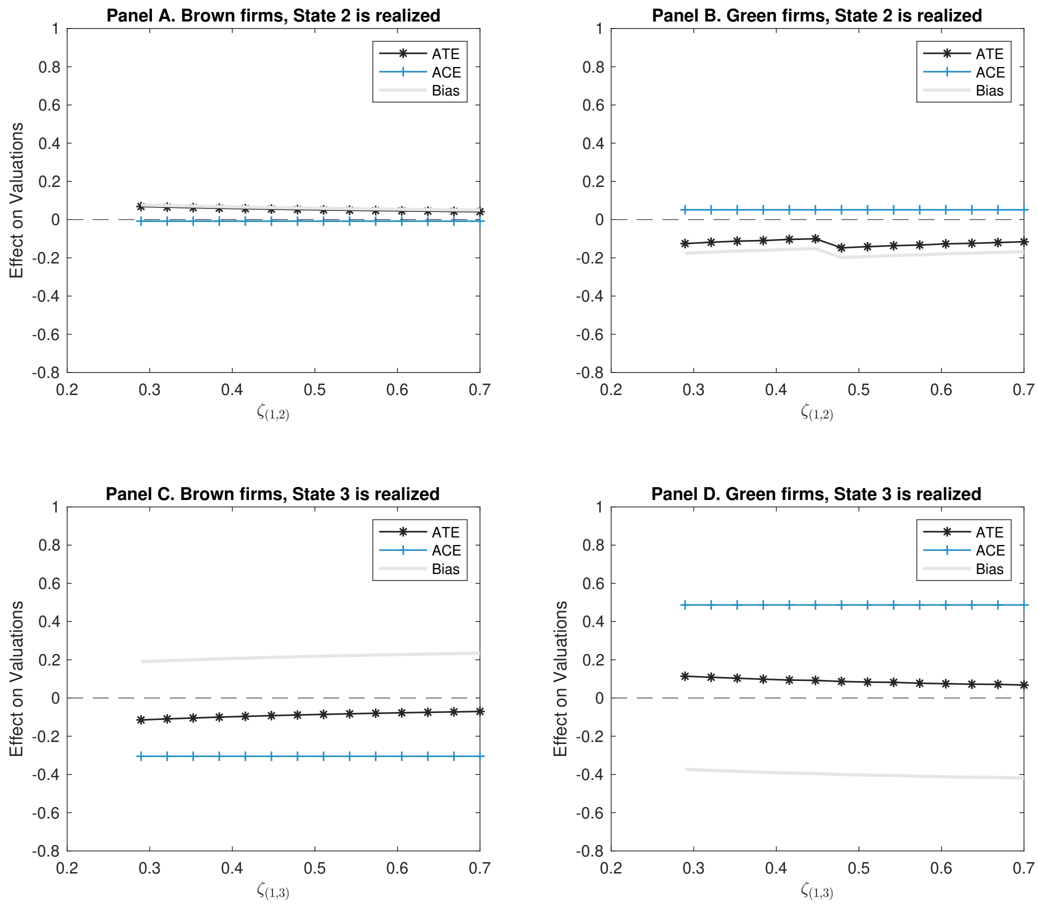

Figure IA.6: Comparative statics: Transition intensities out of State 1. The figures show how the average treatment and causal effects, as well as the bias in treatment effects induced by anticipation vary with the level transition intensities out of State 1. Panels A and B show these comparative statics with respect to the transition intensity from State 1 to State 2, $\zeta_{(1,2)}$, when State 2 is realized. Panel A shows the effects for brown firms, while panel B shows the effects for green firms. Panels C and D show the comparative statics with respect to the transition intensity from State 1 to State 3, $\zeta_{(1,3)}$, when State 3 is realized.

Table IA.1: Baseline calibration. This table shows parameter values in the baseline calibration. Panel A reports the values of the technology parameters, which remain unchanged with the climate regulation state. Panel B shows how the carbon tax rate and equity issuance cost parameters vary across states. Panel C shows the Markov transition matrix of the climate regulation state, whereby each cell represents $\zeta_{(i,j)}$, the transition intensity from state i to state j, with $i, j = \{1, 2, 3\}$.

Panel A. State invariant parameters

Parameter	Symbol	Value
Risk-free rate	r	0.05
Carry cost of cash	λ	0.15
Cashflow growth rate	μ	0.24
Cashflow volatility	σ	0.11
Capital depreciation rate	δ	0.15
Investment adjustment costs	θ	1.8
Correlation between W_t^A and W_t^M	ρ_c	0.4
Market price of cashflow risk	η	0.4
Liquidation value	l	1.00
Correlation with financing risk for green firms	$\rho_{G,\kappa}$	-1.00
Correlation with financing risk for brown firms	$\rho_{G,\kappa}$	1.00

Panel B. Parameters that vary with the state

		State 1	State 2	State 3
Carbon tax rate	$\tau^e_{B,s}$	0	3%	6%
Carbon tax rate	$\tau^e_{G,s}$	0	0	0
Marginal cost of financing for green firms	$\gamma_{G,s}$	6%	6%	6%
Fixed cost of financing for green firms	$\phi_{G,s}$	8%	1%	0.1%
Marginal cost of financing for brown firms	$\gamma_{B,s}$	6%	6%	6%
Fixed cost of financing for brown firms	$\phi_{B,s}$	1%	1.5%	2%
Market price of financing risk	κ_s	1.099	-1.099	-1.099

Panel C. State transition matrix

		State j	
State i	0.4	0.3	0.3
	0	1	0
	0	0	1

Table IA.2: DiD regressions in simulated data. This table reports the regression results from the model in Eq. (IA.2). Specifications (1) and (2) use scaled valuations, $F(m)_{i,t}$ as the dependent variable while specifications (3) and (4) use investment, $i(m)_{i,t}$ as the dependent variable. Standard errors are presented in parentheses. Panel A reports the results when simulating model data that incorporates anticipation (uses the transition intensity matrix in Table IA.1). Panel B reports the results from simulations of model data when assuming no anticipation effects (all diagonal elements of the transition intensity matrix are one, while all off-diagonal elements are zero).

Panel A. Simulations incorporating anticipation

	Valuations		Investment	
Realized state	State 2	State 3	State 2	State 3
	(1)	(2)	(3)	(4)
Treated × *Post*	0.21	-0.14	0.10	-0.13
	(0.0011)	(0.0011)	(0.0001)	(0.0001)
Treated	-0.14	-0.14	-0.06	-0.06
	(0.0011)	(0.0011)	(0.0001)	(0.0001)
Post	-0.14	0.07	-0.07	0.06
	(0.0008)	(0.0008)	(0.0001)	(0.0001)
Intercept	1.14	1.14	0.11	0.11
	(0.0008)	(0.0008)	(0.0001)	(0.0001)
N	200	200	200	200
R-squared	0.71	0.97	0.97	0.99

Panel B. Simulations with no anticipation

	Valuations		Investment	
Treated × *Post*	-0.05	-0.78	-0.04	-0.45
	(0.0009)	(0.0008)	(0.0002)	(0.0003)
Treated	-0.001	-0.001	-0.0001	-0.0001
	(0.0008)	(0.0008)	(0.0002)	(0.0003)
Post	0.04	0.47	0.04	0.29
	(0.0006)	(0.0006)	(0.0001)	(0.0002)
Intercept	1.23	1.23	0.20	0.20
	(0.0006)	(0.0006)	(0.0001)	(0.0002)
N	200	200	200	200
R-squared	0.61	0.95	0.99	0.99

IA.5 General Overview of SMM

The starting point for structural estimation is an economic model which specifies how a firm's choices depend on a variety of fundamental parameters. Usually the values of some parameters are well-documented in existing literature, while the values of other "unknown" parameters are not (e.g., those that cannot be measured empirically). For example, it is straightforward for a researcher to choose a capital depreciation rate from the broad range of existing estimates, but harder to determine a plausible value for brown firms' marginal financing costs in a state with high emissions penalties.

Estimating the economic model using SMM produces an estimated value for each unknown parameter. The researcher starts by making several choices that govern the simulation of the model. First, she pre-sets the values of some fundamental parameters (guided by evidence from prior work), which remain fixed throughout the structural estimation. Second, she designates a set ψ of unknown parameters that are updated after each set of simulations, along with a set of initial guesses $\hat{\psi}$ for the values of the unknown parameters. Third, she decides how to express some of the model's features to facilitate simulation. For example, the researcher can simulate the model either for a single representative firm or for a panel of firms that differ along some characteristic. Moreover, when the model is expressed in continuous time with an infinite horizon, the simulation can be conducted over a fixed number of discrete time periods (e.g., each time period can represent a single day).

Next, the researcher conducts a set of S simulations. In each simulation sim, the researcher draws a random value of the stochastic shock (dW_t in our model) for each firm in the economy. She computes each firm's optimal choices for the first time period, given the values of the shock and the fundamental parameters. Each firm then receives a new random draw for the shock term, and its optimal choices are computed for the second time period (taking into account its choices from period one). The researcher continues this process until the final time period.

After each simulation is finished, the researcher constructs a vector of *simulated moments* $\hat{m}^{sim}(\hat{\psi})$ that are based on the model's outcomes. For example, a moment can be the firm's average choice of cash holdings, or the volatility of its cashflows. The researcher typically chooses an initial "burn-in" period during which the model is simulated, but the resulting outcomes are not used in the moment calculation. This is because model outcomes can vary dramatically across early time periods based on the initial conditions, before eventually stabilizing. The simulated moments are calculated across all time periods after the end of the burn-in period.

Once all S simulations are completed, the researcher calculates the average value of $\hat{m}^{sim}(\hat{\psi})$. Separately, she constructs a vector $\hat{M}$ of *empirical moments* that have the same definitions as the simulated moments, but are calculated using actual data. Given the current (guessed) values in $\hat{\psi}$, the next guess is determined by the value of the objective function :

$$\left(\hat{M} - \frac{1}{S}\sum_{sim=1}^{S} \hat{m}^{sim}(\hat{\psi})\right)' W \left(\hat{M} - \frac{1}{S}\sum_{sim=1}^{S} \hat{m}^{sim}(\hat{\psi})\right), \tag{IA.3}$$

where W is a weighting matrix such as the inverse of the covariance matrix of the empirical moments. These solutions replace the previous values in $\hat{\psi}$. The researcher then proceeds to the second set of S simulations. She repeats this entire process until the values of the simulated moments converge to the values of the empirical moments $\hat{M}$. The values in $\hat{\psi}$ from the final set of simulations are the fundamental parameter values that produce the best fit between the economic model and observed data.

Parameter estimates can be economically interesting in their own right, for example when the corresponding parameters are important for understanding agents' behavior. Moreover, given the estimates, a researcher can conduct a counterfactual analysis of the economic model. She does so by solving the model for (a range of) alternative values for one of the parameters, while keeping all other model parameters fixed. To the extent that the model is a reasonable approximation of reality, the change in the optimal firm outcomes can be used to infer the parameter's causal effect.

One important limitation of structural estimation is that parameter identification become more challenging as the number of elements in ψ increase. Accurate estimation requires at least one moment that is informative for each unknown model parameter. Finding such moments becomes more challenging as the number of unknown parameters grows.