Addressing Anticipation Effects in Finance

Tomislav Ladika Elisa Pazaj Zacharias Sautner

June 9, 2026

Abstract

A wide range of empirical techniques cannot accurately estimate a policy event's causal effects, because agents adjust their 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 the COP 21 meeting, agents assigned a 69% likelihood to an agreement with some form of emissions penalties. Our estimates imply that anticipation led high-emissions firms to reduce investment in the year before the agreement, relative to low-emissions firms. We quantify that reduced-form studies of the Paris Agreement may understate causal effects on investment by 29%, and on stock returns by 22%. Finally, we provide guidance on how our integration procedure can be applied to a wide range of unprecedented but anticipated policy events.

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. We thank Christopher Hennessy, Mathias Kuttli, Arkodipta Sarkar, Enrique Schroth, Lucian Taylor, and Toni Whited, seminar participants at Manchester Business School, and conference participants at the FIRS 2025, EFA 2025, and CFEA 2025 for helpful feedback.

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 an event depend on their beliefs about its possible outcomes, while choices after the event depend on beliefs about further 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 policy change is completely unexpected and permanent (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 lowered investment levels in the data. In both cases, it is challenging for researchers to gauge carmakers' expectations from past policy events, in part because few EV tax credit programs were adopted until 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 (AI) may reshape finance (e.g. Eisfeldt et al., 2026; Babina et al., 2024; Eisfeldt and Schubert, 2026). This latter case highlights how anticipation cannot be fully inferred from pre-trends that are measured shortly before a policy event: Market participants likely formed expectations about possible AI regulations when ChatGPT was released in November 2022, yet any such policy event in the United States is likely still some time away.

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 (henceforth "integration procedure"). 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 procedure 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 climate regulatory risk affect financial outcomes.[1] We

[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., 2026; Seltzer et al., 2026), banks' investments in high-emissions firms (e.g., Alessi et al., 2024), corporate green revenues (e.g., Klausmann et al., 2025), leverage of firms with high climate risk exposures (e.g., Ginglinger and Moreau, 2023), and total factor productivity (e.g., Pang et al., 2023).

show that firms widely expected some form of climate targets, achieved through a penalty on carbon emissions, to be agreed upon at COP 21. As a result, standard reduced-form models significantly underestimate the Paris Agreement's causal effects on firm investment and market valuations.

We first describe our integration procedure in general terms. A researcher writes a dynamic model in which the initial policy state can change over time. This can be 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 actual 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 values 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 one time in the counterfactual absence of anticipation, with transition probabilities set to 0 and other parameters set to the final values. This is analogous to a reduced-form analysis of a natural experiment in which the policy change is exogenous and fully unexpected.

The key novelty is the use of DiD coefficients as moments to match in SMM. These coefficients provide highly useful information for identifying the transition probabilities and state-contingent parameters. By using these coefficients, our integration procedure generates parameter estimates that match theory to 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 final values for model parameters that are estimated *in the presence of anticipation*.

We now describe the application of this procedure 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. Firm beliefs account for the possibility that the state changes multiple times, to account for the policy dynamics discussed by Hennessy and Strebulaev (2020). The firm continuously chooses optimal investment, given its distribution of beliefs about future policy states 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 target, which is then implemented via a specific carbon tax.

The benchmark model produces closed-form expressions that build intuition about the impact of anticipation on firm valuations and policies. Optimal investment decreases in the level of the carbon tax. Both investment and firm value are lower than in an alternative model without the possibility of carbon pricing—even in the initial state without a carbon tax, because valuations and policies incorporate the possibility of future taxation. The benchmark 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.

Climate 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 sensi-

[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, 2026). Our assumption that the carbon tax increases with the capital stock reflects that emissions typically rise with firm size.

tivity of firm value to a carbon tax (Kacperczyk, 2026). 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. Marginal financing costs depend on the policy state and vary for both firms, which in turn choose both investment and precautionary cash savings in anticipation of these changes. Each firm's cost of capital adjusts to the resulting competitive advantage for low-emissions firms.[3] This setup can represent a government that implements the Paris Agreement by levying a carbon tax on high-emissions 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. Our model contributes to other pioneering theories of the dynamic effects of carbon emissions regulation on firms (e.g., Bustamante and Zucchi, 2024; Albuquerque et al., 2026), 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*, *Moderate* or *Strict* state, which differ in the carbon tax rate on high-emissions firms. Tax rates are calibrated using Social Cost of Carbon (SCC) estimates from Nordhaus (2019), with the *Strict* state corresponding to a goal of limiting temperature increases to 1.5℃ above pre-industrial levels. Conditional on carbon taxes being adopted, our model allows for the emissions penalties to be either incrementally strengthened or watered down in future time periods.

Model simulation can produce multiple types of state changes, but empirically we only observe the Paris Agreement as the outcome of COP 21. In order to match model-simulated outcomes to actual data, we need to map the agreement to one of the model's policy states, and we choose the *Strict* state. One reason is that contemporary observers widely regarded the agreement as more ambitious than anticipated. Financial media described the Paris Agreement as "delivering a more ambitious goal than had been expected... The agreement

[3] The importance of financial constraints for understanding the effects of carbon pricing is documented, for example, in Döttling and Rola-Janicka (2025).

surpassed what had been anticipated" (The Economist, 2015). Also, the Paris Agreement explicitly adopted a goal of limiting warming to 1.5°C, which was the most ambitious target under serious consideration at COP 21.[4] Further, model estimation validates this choice: SMM with the Paris Agreement mapped to the *Mild* or *Moderate* state consistently fails to match the empirical DiD moments.

Next, we estimate DiD regressions following Bolton and Kacperczyk (2021) (henceforth "BK"), using both model-simulated and empirical data. BK examine how stock returns change for high- versus low-emissions firms, in the year before versus after the Paris Agreement. We also examine investment as an additional outcome variable. Our simulated DiD coefficients very closely match those obtained using real data, resulting in highly precise estimates for most transition intensities. Our process also produces a reasonable match for 16 additional moments, allowing us to estimate other unknown parameters such as the financing costs in each state.

Our results indicate that there was significant anticipation that a carbon tax would be adopted at COP 21, but that it would most likely be modest in size. We estimate that prior to the meeting, firms assigned a 59.5% probability to the *Mild* state, a 9.8% probability to the *Moderate* state, and a 0.03% probability to the *Strict* state (implying a 31% likelihood that no carbon tax would be adopted). In other words, firms' observed responses to the Paris Agreement are only consistent with our model if the level of anticipation of some type of emissions penalty was high, and the expected tax level was modest. Nevertheless, high-emissions firms' investment rates declined significantly after the Paris Agreement relative to low-emissions firms, consistent with the realized agreement being stricter than expected.

Our counterfactual analysis solves the model for a completely unanticipated switch from the *Pre* to *Strict* state, and compares the resulting DiD regression results to the empirical DiD estimate. Our results indicate that DiD models which do not account for anticipation underestimate the Paris Agreement's causal effect on investment by 29%, and on stock returns by 22%. This implies that before the Paris Agreement was announced, brown firms had already reduced investment in anticipation of a regulatory change, and their market

[4]Much doubt has been cast on whether the Paris signatories will fulfill their pledges, especially after the United States pulled out of the agreement. Our model incorporates this uncertainty, by allowing firms to form beliefs that the stated climate targets would be reduced over time.

valuations already reflected substantial parts of the impact of carbon taxes.

Finally, we propose two ideas for generalizing our procedure to study other types of unprecedented but anticipated policy events, which require adding only a few steps to a standard reduced-form analysis. First, the carbon tax parameter in our model can represent numerous other firm frictions or costs, allowing researchers to study how investment responds to tariffs (e.g., the "Liberation Day" announcements in the U.S.) or technological shocks that reduce operating costs (e.g., the release of AI agents). Second, our integration procedure can apply to any observed firm outcome that is informative about anticipation of policy changes. One candidate outcome is changes in corporate political spending, which could be used to produce estimates of firms' beliefs about many types of policy events.

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), to rule out a violation of the compound exclusion restriction when re-using a natural experiment (Cronqvist et al., 2024), or to quantify reduced-form evidence (Catherine et al., 2022). 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 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., 2026), 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 and external financing costs that vary by firm type and policy state. This full model provides economic structure for a DiD comparison of high- and low-emissions 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 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 a state before any major climate regulations are adopted, and is analogous to the *Pre* period of a DiD model analyzing the Paris Agreement. Carbon taxes are positive in the other $N - 1$ states, with $0 = \tau_1^e < \tau_2^e < \cdots < \tau_N^e$.

Given an economy in state s at 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.[5]

Dynamics of capital and profits. At each time t the firm invests in new capital, earns profits, and pays out dividends. The firm's capital stock dynamics 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 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 (defined below) and $\Gamma(I_t, K_t)$ 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. $g(i_t)$ is increasing and convex in i_t and 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. Eq. (3) is often used in the literature to represent a firm that incurs

[5] 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).

adjustment costs only when net investment is non-zero.

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 the policy states, because the model's focus is on studying how firm outcomes vary with emissions penalties (and external financing costs in Section 2.3).[6]

Firm value. Firm value equals the expectation of all discounted proceeds to shareholders until liquidation:

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

where the expectation is taken under the risk-neutral measure $\mathbb{Q}$, and $e^{-\int_0^t r_u du}$ 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_A. The risk-adjusted cashflow shock under the risk-neutral measure $\mathbb{Q}$ is:

$$dA_t = \hat{\mu}_A dt + \sigma_A d\hat{W}_t^A, \tag{5}$$

where $\hat{W}_t^A$ is a standard Brownian motion under the risk neutral measure $\mathbb{Q}$ and $\hat{\mu}_A$ is the risk-adjusted mean cashflow shock.[7] The following risk-adjustment applies:

$$\hat{\mu}_A = \mu_A - \rho_A \sigma \eta_A, \tag{6}$$

where ρ_A is the correlation between firm and aggregate cashflows, and η_A the market price of cashflow risk.

[6]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.

[7]To highlight the key mechanisms, we initially model cashflows as an arithmetic Brownian motion, which implies that all shocks are transitory. In Section 2.3, we relax this assumption by introducing persistence, in line with extensive empirical evidence of persistent firm-level cashflows (e.g., Hennessy and Whited, 2007; Gryglewicz et al., 2022).

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:

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

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 $\hat{\mu}_A K$, net investment $(i - \frac{\theta(i-\nu)^2}{2})K$, 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:

$$q_s = 1 + \theta(i_s - \nu), \tag{8}$$

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:

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

where q_{s^-} is the marginal q in each other state s^-. Average 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 $\hat{\mu}_A > \tau_s^e + r + \delta + \sum_{s \neq s^-} \zeta_{s,s^-}(1 - q_{s^-})$. 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 state and on the marginal q_{s^-} values therein. 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 derived from Eq. (9) is:

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

This expression represents the causal effect of a tax policy on investment in the presence of anticipation, and it is always negative since $\theta > 0$. This means that an increase in the 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). Note that this result differs from Albuquerque et al. (2026), who find that high-emissions firms may optimally *in-*

crease investment prior to a carbon tax adoption. 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 state change can occur any time in our model, compared to just in 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:

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

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 $\hat{\zeta}_{s,s^-} \equiv e^{\kappa_s}\zeta_{s,s^-}$ is 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." To understand the intuition for this premium, first consider a case in which firm value is higher in state s than other states s^-, that is, 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, that is, climate regulation boosts firm value. The term is negative overall, implying that anticipation of a potential switch to better states decreases the transition risk

[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.

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, introduce persistence in cashflows, 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 climate regulation, as well as market participants' prior expectations that any such 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 (2026), 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 marginal cost of $\gamma_{f,s}$. State-contingent financing costs represent the possibility that in a state with high emissions penalties, brown firms face more difficulty raising equity financing (e.g., due to higher information 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 funds). We denote the cumulative external financing process by E, with dE_t denoting financing raised in each time increment.

Cashflows. The risk-adjusted cashflow shock scales with the firm's underlying productivity P_t. Under the risk-neutral measure $\mathbb{Q}$ cashflow shocks are given by:

$$dA_t = \hat{\mu}_A P_t dt + \sigma_A P_t d\hat{W}^A_t, \tag{12}$$

where $\hat{\mu}_A$ is the risk-adjusted drift of the productivity-scaled cashflow process, σ_A is the diffusion term and $\hat{W}_t^A$ is a standard Brownian motion under the measure $\mathbb{Q}$.

Shocks to firm productivity are permanent, with productivity P_t given by:

$$\frac{dP_t}{P_t} = \hat{\mu}_P dt + \sigma_P d\hat{W}_t^P, \tag{13}$$

where $\hat{\mu}_P$ is the risk-adjusted drift of the productivity process, σ_P is the diffusion term and $\hat{W}_t^P$ is another standard Brownian motion under the risk-neutral measure $\mathbb{Q}$. The two Brownian motion terms are correlated: $d\hat{W}_t^P d\hat{W}_t^A = \rho dt$. The following risk-adjustment applies:

$$\hat{\mu}_P = \mu_P - \rho_P \sigma_P \eta_P, \tag{14}$$

where ρ_P is the correlation between firm and aggregate productivity shocks, and η_P is the market price of productivity risk. The risk adjustment for the drift of the productivity-scaled cashflow process $\hat{\mu}_A$ is still as given by Equation (6).

The firm's incremental operating profit dY_t now scales with capital K_t and productivity P_t, and is determined by the following:

$$dY_t = dA_t K_t - I_t P_t dt - \Gamma(I_t, K_t) P_t dt - \tau_s^e K_t P_t dt, \tag{15}$$

where $\Gamma(I_t, K_t)$ is defined as in the benchmark model.

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:

$$dM_t = dY_t + (r - \lambda) M_t dt - dD_t + dE_t, \tag{16}$$

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:

$$\mathbb{E}^{\mathbb{Q}}\left[\int_0^{\tau} e^{-rt}(dD_t - dE_t - \gamma_{f,s}dE_t) + e^{-r\tau}(\omega K_\tau P_\tau + M_\tau)\right] \tag{17}$$

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. (17) shows that in each state, firm value depends on the cash stock M_t, capital stock K_t, and the productivity level P_t. Let $V(M,K,P,s)$ 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 firm 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 $M \in (\underline{M}_s, \overline{M}_s)$, firm value satisfies the following HJB equation under the risk-neutral measure:

$$\begin{aligned} rV(M,K,P,s) = \max_{i_s} \; \Bigg\{ & \left[\hat{\mu}_A KP - i_s KP - \frac{\theta(i_s-\nu)^2}{2} KP - \tau^e_{f,s} KP + (r-\lambda)M\right] V_M(M,K,P,s) \\ & + \frac{1}{2}K^2P^2\,\sigma_A^2 V_{MM}(M,K,P,s) + V_K(M,K,s)(i_s-\delta)K \\ & + \hat{\mu}_P V_P(M,K,P,s) + \frac{1}{2}\sigma_P^2 V_{PP}(M,K,P,s) \\ & + \sum_{s=1}^{N-1} \hat{\zeta}_{s,s^-}\left[V(M,K,P,s^-) - V(M,K,P,s)\right] \Bigg\} \end{aligned} \tag{18}$$

Similar to the HJB (7) of the benchmark model, the RHS of Eq. (18) details how each of the state variables (M,K,P) 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

[9] When cash holdings become perilously low, the firm can either liquidate or raise significant external financing. The threshold $\underline{M}_s$ arises when liquidation provides greater firm value than continuation after raising costly financing.

growth rate $\hat{\mu}KP$, total investment and associated adjustment costs $i_sK+\theta(i_s-\nu)^2KP/2$, the carbon tax $\tau^e_{f,s}KP$, 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 following two terms capture the effects of the changes in the productivity process P_t.

The final term in Eq. (18) 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, $V(M,K,P,s^-) - V(M,K,P,s)$. Unlike in the benchmark model, this sensitivity of firm value to the state change now depends on the firm's cash holdings and physical capital. This can represent that firm value decreases by more after the introduction of carbon taxes for firms with less precautionary savings, as they 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 cash, capital and productivity:

$$V(M,K,P,s) = PKV(\frac{M}{PK},s) \equiv PK\ F(m,s), \tag{19}$$

where $m \equiv M/PK$ 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 penalties on firms' investment decisions, taking into account the anticipation of changes in climate regulation states. Note also that $V_K = P(F(m,s) - mF'(m,s))$, $V_M = F'(m,s)$, $V_{MM} = \frac{1}{PK}F''(m,s)$, $V_P = K(F(m,s) - mF'(m,s))$, $V_{PP} = K\frac{m^2}{P}F''(m)$ and $V_{PM} = -\frac{m}{P}F''(m)$.

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

$$
\begin{aligned}
rF(m,s) = &\left[\hat{\mu}_A - i_s - \frac{\theta(i_s-\nu)^2}{2} - \tau^e_{f,s} + (r-\lambda)m\right] F'(m,s) + \frac{1}{2}\sigma_A^2 F''(m,s) \\
&+ \hat{\mu}_P \left[F(m,s) - mF'(m,s)\right] + \frac{1}{2}\sigma_P^2 F''(m,s) - m\sigma_P\sigma_A\rho F''(m) \\
&+ (i_s - \delta)\left[F(m,s) - mF'(m,s)\right] + \sum_{s=1}^{N-1} \hat{\zeta}_{s,s^-} \left[F(m,s^-) - F(m,s)\right]
\end{aligned}
$$

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

$$
\begin{aligned}
(r - \hat{\mu}_P - i_s + \delta)F(m,s) = &\left[\hat{\mu}_A - i_s - \frac{\theta i_s^2}{2} - \tau^e_{f,s} + (r - \lambda - \hat{\mu}_P - i_s + \delta)m\right] F'(m,s) \\
&+ \frac{1}{2}(m^2\sigma_P^2 + \sigma_A^2 - 2m\sigma_A\sigma_P\rho)F''(m,s) \\
&+ \sum_{s=1}^{N-1} \hat{\zeta}_{s,s^-} \left[F(m,s^-) - F(m,s)\right]
\end{aligned}
\tag{20}
$$

We derive the firm's first-best choices by solving Eq. (20) numerically, using the upper and lower boundary conditions on cash m for each state s. The first-order condition for investment is given by:

$$
i^*(m,s) = \frac{1}{\theta}\left(\frac{F(m,s)}{F'(m,s)} - m - 1\right) + \nu \tag{21}
$$

Eq. (21) 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 integration procedure adds three features to a standard structural estimation using SMM.[10] First, the model allows for one or more state changes, 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, a process that is determined by a matrix of transition intensities 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 two possible states (s_1, s_2), the transition matrix is:

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

where the element in row i and column j is the transition intensity from state i to state j.

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 the stochastic process determined by the transition intensity values in $\hat{\psi}$. Because the policy state applies to the aggregate economy, when a state change occurs, it affects 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.

[10] Internet Appendix IA.2 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.

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 $\zeta_{s_2,s_2} = 1$ (and thus $\zeta_{s_2,s_1} = 0$) implies that the economy remains in state s_2 permanently once it arises.

3.2 Using DiD coefficients as moments

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 and a state change to mimic the policy event.[11] Also, a state change's impact on the parameters in ψ should differ across firm types. For example, in our application, 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 Internet Appendix IA.2).[12]

The moments obtained from DiD regressions are highly informative for identifying unknown state-contingent parameters (including the transition intensities), because changes to

[11] DiD regressions can also 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 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.

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.

Model simulation can produce state changes that have not occurred in real life. Consider a researcher who uses a three-state Markov chain and initially chooses positive 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 world market participants have anticipated but not experienced.

Our procedure needs to map the observed post-shock state to a specific state 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, 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}.

3.3 Comparison with calibration

Hennessy and Strebulaev (2020) and Hennessy and Livdan (2021) provide important contributions to help researchers account for anticipation effects during model calibration. In contrast, our integration procedure focuses on model estimation.[14] A key reason is that treated

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

[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.

firms in most DiD analyses are usually different from the average sample firm, but existing literature often provides estimates of economic model parameters only for the latter. As such, researchers have little guidance on how to pre-set the values even of standard parameters for green and brown firms. This limitation is evident in the firm value expression in Eq. (17), which depends on firm- and state-specific financing costs that cannot be readily calibrated.

Model estimation can provide several additional benefits over calibration. It can yield statistics regarding the uncertainty of parameter estimates. Additionally, 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. 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 not take into account important channels.

4 Economic model estimation

This section applies our integration procedure to the full economic model developed in Section 2.3. We describe the sample and data used to calculate empirical moments and the process for estimating or pre-setting model parameters. We report the results of the model's estimation, and also a counterfactual analysis that quantifies the anticipation bias in estimates of the Paris Agreement's causal 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 sampling procedure begins with all U.S. firms in the intersection of these databases. To identify treated firms (corresponding to brown firms in our model), we first measure each firm's emissions intensity as Scope 1 emissions scaled by total revenue. We measure this ratio in 2014, to account for the typical lag in reporting of emissions data identified by Zhang (2025). Scope 1 emissions are 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). We classify a firm as treated if its emissions intensity is above the 66^{th} percentile in 2014. To identify control firms (corresponding to green firms in our model), we follow BK and match each treated firm to one firm with an emissions intensity below the 66^{th} percentile, using the nearest neighbor methodology without replacement.[15]

This procedure results in a sample of 622 firms. The sample covers the last three months (quarter Q4) of 2014 through Q4 of 2016. We omit Q4 of 2015, since it is not clear whether firm outcomes in that quarter were determined before or after the Paris Agreement's announcement on December 12, 2015. All empirical moments are measured using data from this sample. Table 1 presents detailed definitions and summary statistics for the variables used to construct the moments, including the DiD regression variables in Section 4.3.

4.2 Model parameters

4.2.1 Parameters to be estimated

The transition intensities, state-contingent financing costs, and investment adjustment costs are difficult to measure outside the model, and there is little evidence in existing literature on their values—especially for the green and brown firms in our setting. Hence, we cannot reliably pre-set these parameters' values and instead estimate them.

We assume there are four possible climate regulatory states. The starting *Pre* state involves no regulation, which implies no taxes on carbon emissions ($\tau^e_{Pre} = 0$) and no differences in financing costs between green and brown firms ($\gamma_{G,Pre} = \gamma_{B,Pre}$). The other three states involve some climate regulation, with positive taxes on carbon emissions ($\tau^e_s > 0$) and possible differences in external financing costs between green and brown firms. We denote these states as *Mild*, *Moderate* and *Strict*, with $\tau^e_{Mild} < \tau^e_{Moderate} < \tau^e_{Strict}$.

The Markov chain that governs transitions between the policy states is expressed as a 4×4 matrix, whose rows/columns correspond to the *Pre*, *Mild*, *Moderate* and *Strict* states:

[15] We construct propensity scores by matching on the same set of variables as BK. We match on all variables listed in Table 1, except for *Stock return* and *Equity issuance*.

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

where each off-diagonal element $\zeta_{s,s'}$ denotes the transition intensity from state s to state s'.[16]

This transition matrix captures the dynamic nature of climate policies, by allowing for regulation to either tighten ($\zeta_{Moderate,Strict} \geq 0$) or weaken after initial adoption ($\zeta_{Moderate,Mild} \geq 0$ and $\zeta_{Strict,Moderate} \geq 0$). We make two key assumptions that reduce the number of additional transitions. First, once some climate regulation has been adopted, it is no longer possible to return to a state without regulation ($\zeta_{Mild,Pre} = \zeta_{Moderate,Pre} = \zeta_{Strict,Pre} = 0$). Second, additional transitions are only possible between adjacent states ($\zeta_{Mild,Strict} = 0$ and $\zeta_{Strict,Mild} = 0$). These assumptions reflect the political reality that regulations are rarely repealed in entirety, and also tend to be adjusted gradually.[17] In our setup, the transition intensities do not change over time. This implies that in the Pre period, firms already form conditional beliefs about the likelihood that different regulations will be adjusted after adoption.

We estimate green and brown firms' financing costs in the *Strict* state, as well as the common financing cost in the Pre state. We further assume that financing costs change linearly with the degree of environmental regulation. This simplifies the estimation, by allowing us to directly calculate green and brown firms' costs in the *Mild* and *Moderate* states by linearly interpolating between their respective costs in the *Pre* and *Strict* states.

Overall, our set ψ contains 12 unknown parameters to estimate:

$$
\psi = \{\ \zeta_{Pre,Mild},\ \zeta_{Pre,Med},\ \zeta_{Pre,Strict},\ \zeta_{Mild,Med},\ \zeta_{Med,Mild},\ \zeta_{Med,Strict},\ \zeta_{Strict,Med},\ \theta_G,\\
\theta_B,\ \gamma_{Pre},\ \gamma_{G,Strict},\ \gamma_{B,Strict}\}.
$$

We estimate the full model using SMM to obtain a value for each parameter in ψ. We

[16] Transition intensities are not probabilities, and need not have an upper bound of 1, but there is a direct mapping between intensities and probabilities. Specifically, transition probabilities over a discrete time interval dt are $P(dt) = e^{\mathbf{Z}\cdot dt}$, where $\mathbf{Z}$ is the intensity matrix.

[17] Our procedure can be easily extended to allow for additional transitions between the states. For example, allowing mild emissions penalties to later be abandoned only requires estimation of one additional parameter $\zeta_{Mild,Pre}$.

simulate an economy of 490 firms (a similar order of magnitude to our empirical data), whose optimal choices are computed at a quarterly 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. 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, estimated using the influence function approach (Erickson and Whited, 2000). Table 5 lists all of the moments used to identify the parameters in ψ.

Following Section 3.2, we map the Paris Agreement to the *Strict* state during estimation, and designate the *Mild* and *Moderate* states as counterfactuals that did not arise. One reason for this choice is that the Paris signatories agreed to a goal of limiting temperature increases to 1.5°C, which requires emissions penalties in line with high 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* or *Moderate* states arise or no state change occurs. Notably, simulated model outcomes in the *Pre* state still depend in part on the transition intensities to the *Mild* or *Moderate* states and the financing costs in these states. This is because firms' optimal choices prior to the adoption of climate regulation take into account the possibility that these states might arise (and also that the *Moderate* state is reversible).

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 is $\rho = 0.4$. We set the liquidation value to $\omega = 1$.

We calibrate the carbon tax rates to $\tau_{B,Mild} = 0.49\%$, $\tau_{B,Moderate} = 1.83\%$ and $\tau_{B,Strict} = 3.16\%$ based on SCC estimates in Nordhaus (2019). Specifically, we first compute the average Scope 1 emissions per dollar of assets across all treated firms in our sample. To obtain $\tau_{B,Mild}$ we multiply this amount by \$37, which is the less stringent policy target estimate for 2015. Similarly, $\tau_{B,Strict}$ equals average emissions multiplied by \$236, which represents the most ambitious target of limiting temperature increases to 1.5℃. We set $\tau_{B,Moderate}$ as the midpoint between the *Mild* and *Strict* tax rates.

We set the κ_s parameters to $\kappa_{Pre} = ln(3)$, $\kappa_{Mild} = -ln(3)$, $\kappa_{Moderate} = -ln(3)$ and $\kappa_{Strict} = -ln(3)$. These values are based on BCW, and they represent the fact 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 first match firms across our and their samples, with a successful match for 348 out of 622 firms. For the matched firms we then compute the median cashflow growth rate and volatility, productivity growth rate and volatility, and the correlation between cashflow shocks and productivity shocks.

4.3 Identification of unknown model parameters

4.3.1 Transition intensities

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

$$Y_{i,t} = \beta_0 + \beta_1 \times Treated \times Post + \beta_2 \times Treated + \beta_3 \times Post + \gamma X_{i,t} + \psi_i + \epsilon_{i,t}, \quad (22)$$

We estimate Eq. (22) using actual and simulated data for the quarterly outcome variables $Y_{i,t}$ of *Stock return* and *Investment*.[18] In the actual data, $Treated$ equals 1 for high-emissions firms and 0 for matched low-emissions firms, as defined in Section 4.1. $Post$ equals 1 for Q1 through Q4 of 2016, and 0 for Q4 of 2014 through Q3 of 2015. In the simulated data, we designate half of firms in the economy as brown and half as green. $Post$ 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 (i.e., during the initial *Pre* period).

[18] Our theoretical model does not include debt, so stock returns in the simulated dataset are unlevered by definition. To facilitate proper comparison of empirical and simulated moments, we also un-lever *Stock return* when measured using actual data.

Table 2 shows regression results estimated on actual data. In columns (1) and (2), the positive coefficient indicates that after the Paris Agreement, stock returns of brown firms rose relative to those of green firms. This is consistent with the evidence in BK. In columns (3) through (5), the negative interaction term coefficient implies that brown firms significantly cut investment relative to green firms after the agreement. The coefficient estimates are statistically significant at the 1% level in all specifications, and robust to inclusion of firm fixed effects and various controls based on BK.[19]

BK interpret the relative increase in brown firms' stock returns 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, in our theoretical model the market price of risk in Eq. (12) is not state-contingent. Instead, the relative increase in brown firms' stock returns arises because the *Pre* period transition risk premium was large enough that, even after accounting for the capital loss from moving to *Strict*, the total realized return for brown relative to green in the post-period exceeds the brown-green return differential in the *Pre* period.

When estimating the model parameters in ψ, we compare the interaction term coefficients obtained from using actual and simulated data, when estimating Eq. (22) based on the models in columns (1) and (4). We use these regression models because in the simulated data we cannot compute most control variables, as they have no equivalents in the theoretical model. The moment values are reported as $\beta_{1,Inv}$ and $\beta_{1,Ret}$ in Table 5.

Our integration procedure requires that the DiD coefficients are informative for identifying the transition intensities in ψ. We verify this 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 theoretical model repeatedly by changing only the transition intensities, and re-estimate Eq. (22) using the resulting simulated dataset. If the DiD results change with the transition intensities, then this is strong evidence that the regression-based moments provide unique identifying information.[20]

[19] Table IA.3 documents similar results when classifying treated firms based on the 75^{th} percentile of the Scope 1 emission intensity in 2014, and then re-matching each of these firms to a control firm with emissions below the 75^{th} percentile.

[20] In these comparative statics, as we vary the transition intensity from the *Pre* state to another regulatory state $s^{'}$, we keep the total exit intensity $(\zeta_{Pre,Mild}+\zeta_{Pre,Moderate}+\zeta_{Pre,Strict})$ constant. We do so to examine only the impact of increasing the probability that, conditional on leaving state *Pre*, the economy moves directly into state $s^{'}$. This probability is effectively computed as $\zeta_{Pre,s'}/(\zeta_{Pre,Mild}+\zeta_{Pre,Moderate}+\zeta_{Pre,Strict})$,

Figure 1 presents the estimates separately for the interaction-term coefficients $\beta_{1,Inv}$ and $\beta_{1,Ret}$, and for transitions from the *Pre* state to either the *Mild* or *Strict* states. Both simulated DiD coefficients vary monotonically with the transition probabilities. Importantly for identification, increases in the transition intensities to the *Mild* and *Strict* states have opposite effects on the coefficient estimates, which is consistent with how anticipation should affect firm outcomes. Specifically, when brown firms anticipate that the future carbon tax is most likely to be small (corresponding to a high intensity $\phi_{Pre,Mild}$), a subsequent realization of the *Strict* state comes as more of a surprise, and leads brown firms to reduce investment by a large amount (relative to green firms). The opposite is true when brown firms expect a future carbon tax to most likely be high (a high intensity $\phi_{Pre,Strict}$). They begin to cut investment during the *Pre* state, and a realization of the *Strict* state then leads to a small additional reduction in investment. For similar reasons, brown firms' valuations of increase more when the realization of the *Strict* state comes as more of a surprise (i.e., when $\phi_{Pre,Mild}$ is higher or $\phi_{Pre,Strict}$ is lower).

4.3.2 Investment adjustment costs

The primary moments used to identify θ_G and θ_B are the variances of investment for green and brown firms over the whole sample period. The variance of investment is an informative moment because firms smooth investment more when adjustment costs are larger. The interaction term coefficient from Eq. (22), 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 four moments that capture the covariance between investment and cash holdings conditioning on firm type and time relative to the adoption of the carbon tax.

and is denoted as $\phi_{Pre,s'}$ in Figure 1.

4.3.3 Financing costs

To help in the identification of the unknown financing cost parameters we also use moments that relate to the frequency, size and variance of equity issuance. The model allows the financing costs to vary across states and firm types, so we construct four separate moments for each statistic conditioning on green and brown firms, before and after the adoption of a carbon tax.

4.4 Structural estimation results

Table 3, Panel A, presents estimates of each unknown parameter in ψ. The first seven rows report estimates of each transition intensity, with the transition probabilities obtained from these values reported in Table 4. Rows 1 through 3 of Table 3 show that the highest intensity from the *Pre* state is $\zeta_{Pre,Mild} = 0.805$, corresponding to a probability of 0.595. The intensity from the *Pre* to *Moderate* state is roughly half as large, corresponding to a 0.098 probability. The intensity to the *Strict* state is almost 0. The three transition intensities jointly imply a 0.307 probability of remaining in the *Pre* period. Overall, the results indicate that firms' observed choices around the COP 21 meeting are consistent with anticipation of a relatively low carbon tax as the most likely meeting outcome. The data also indicate that the strict level of climate goals that are actually contained in the Paris Agreement were almost completely unanticipated by firms.

As shown in rows 4 through 7 of Table 3, firms further anticipated that if higher carbon taxes were to be adopted at COP 21, then they would more likely be rolled back at some future time than to be increased. Conditional on a switch to the *Moderate* state, firms anticipated carbon taxes would later be reduced with 77.9% probability, remain constant with 22% probability, and increase with just 0.1% probability (the intensity $\zeta_{Moderate,Mild} = 1.514$ is much higher than $\zeta_{Moderate,Strict} = 0.003$). Conditional on a switch to the *Strict* state, firms anticipated regulation would later be loosened to the *Mild* or *Moderate* level with a combined 30.5% probability. The transition intensity $\zeta_{Strict,Moderate} = 0.365$ implies that firms anticipated the highest carbon tax rate (if adopted) would remain in effect for about 33 months before being reduced.[21] Interestingly, before COP 21 firms anticipated that if the

[21]Standard discounted cashflow reasoning suggests that near-term cashflows account for only a small proportion of firm value, but this logic assumes that the terminal value is unaffected by the temporary

$Mild$ state arose, it would essentially become an absorbing state. This is consistent with firms anticipating that a low carbon tax regime would almost certainly not be expanded over time.

The standard errors on the transition intensity estimates show that $\zeta_{Pre,Mild}$, $\zeta_{Pre,Moderate}$, $\zeta_{Moderate,Mild}$, and $\zeta_{Strict,Moderate}$ are measured precisely, with t-statistics well above 2. The transition intensity $\zeta_{Pre,Strict}$ is measured less precisely, as are the additional transition intensities to higher carbon taxes $\zeta_{Mild,Moderate}$ and $\zeta_{Moderate,Strict}$. This is likely because each of these intensity values is close to its lower bound of 0.

The near-zero transition intensities reflect a joint identification requirement imposed by two large and opposing data patterns. The data require a high *Pre* period transition risk premium for brown firms so that its collapse at the Paris announcement generates a positive return DiD, while simultaneously requiring that the *Strict* outcome was not priced in so that investment falls sharply. Any direct probability assigned to $\zeta_{Pre,Strict}$ destroys both responses by compressing the gap between pre- and post-Paris valuations. The escalation intensities $\zeta_{Mild,Moderate}$ and $\zeta_{Moderate,Strict}$ are driven to zero by the same logic, though their impact is indirect.

The investment adjustment cost parameters $\theta_G = 18.29$ and $\theta_B = 62.24$ are both large and statistically significant. The higher adjustment costs for brown than green firms reflect their greater difficulty to restructure capital stock in response to stricter climate regulation. This is consistent with higher costs of retiring or adapting existing brown assets. Finally, we estimate that before COP 21 the equity issuance cost was $\gamma_{Pre} = 0.187$, and this cost fell sharply for green firms after the Paris Agreement to $\gamma_{G,Strict} = 0.035$. This is consistent with investors rewarding these firms with cheaper equity access in a strict environmental regulatory state. For brown firms, the issuance cost more than doubled to $\gamma_{B,Strict} = 0.458$.

Table 5 lists the 18 moments that the SMM attempts to match during model estimation,

regulatory spell. In our model, firm value satisfies $V(M, P, K, s) = PKF(m, s)$, which means that the terminal value is directly proportional to the capital stock K. Since productivity P follows a permanent process, investment suppression during the *Strict* regulatory spell permanently reduces K relative to the no-Paris counterfactual, and because all future cashflows scale with KP, the entire future earnings stream—including the terminal value—is affected. The financial frictions channel further amplifies this effect. Higher financing costs for brown firms in the *Strict* state lead to reductions in investment beyond what the carbon tax alone would imply. By the same reasoning, lower financing costs for green firms significantly increase their investment and valuations. It is therefore the interaction between permanent productivity, capital accumulation and state-contingent financing costs that allow a 33-month strict spell to generate the significant investment and return DiD moments observed in the data.

and reports their simulated and empirical values. The DiD coefficient estimates obtained using model-simulated data very closely match the corresponding estimates obtained using actual data, for both stock returns and investment. These moments carry the most direct identifying information for the transition probabilities and the causal effects that are the focus of our analysis. The precision of the match is important for our counterfactual analysis: As discussed in Section 4.5, anticipation bias is measured as the difference between the empirical coefficient estimates in Table 5 and simulated coefficients obtained from estimating the model without anticipation. This counterfactual exercise is less informative when the simulated DiD coefficients with anticipation do not fit the actual data well.

The model does reasonably well at matching the equity issuance moments in the Pre period. Simulated issuance frequencies and amounts for both green and brown firms are all close to their empirical counterparts. However, in the model green firms optimally issue more equity following the adoption of high carbon taxes than in the actual data, with a simulated issuance frequency value of 0.046 compared to an empirical value of 0.009. One possible explanation is that green firms in reality faced fewer investment opportunities in the year after the Paris Agreement, which our model cannot capture because investment opportunities do not change with environmental regulatory state.

The model fits the covariance moments less precisely, though this is of limited concern since both the simulated and empirical values are close to 0. The model correctly captures the key qualitative feature of the data—that cash holdings and investment are weakly related within firms around the Paris Agreement—even if the small quantitative discrepancies are not fully resolved.

4.5 Counterfactual analysis

To gauge the magnitude of how anticipation affects firm outcomes around the Paris Agreement, we solve for firms' optimal responses to the adoption of high carbon tax rates in the counterfactual absence of any anticipation. Specifically, we run a set of model simulations with all transition intensities except $\zeta_{Pre,Pre}$ set to 0, so firms assign a probability of 1 to staying in a regulatory state without carbon taxes. We set all unknown model parameters in ψ to the values estimated in Table 3. In this counterfactual, firms make optimal choices without considering carbon tax rates or financing costs outside their current state (i.e. each state is considered permanent). The model begins at $t = 1$ in the *Pre* state, and we then artificially impose a change to the *Strict* state in a future time period (necessary since $\zeta_{Pre,Strict} = 0$ in the counterfactual Markov chain matrix). We then re-estimate Eq. (22) using the simulated dataset obtained from this set of counterfactual simulations. This exercise is analogous to a reduced-form analysis of an exogenous policy change that is completely unanticipated, and it provides quantitative estimates of the causal effects of the Paris Agreement.

Table 6 compares the DiD coefficients from the counterfactual simulation to the empirical estimates. The counterfactual estimates represent the true causal effects of the Paris Agreement, since they are not affected by anticipation. The counterfactual DiD estimate $\beta_{1,Inv}$ is -0.0021 (standard error of 0.0001), while the empirical estimate is -0.0015 (standard error of 0.0004). This implies that anticipation biases the true effect of the Paris Agreement on investment by 29% (standard error of 0.03). The economic interpretation is straightforward: Because brown firms anticipated the adoption of relatively low carbon taxes at COP 21, they already invested less before the meeting (relative to green firms) than they optimally would have in the absence of anticipation. The realization of ambitious climate goals in the Paris Agreement therefore produced smaller investment changes than would have occurred, had firms completely not expected any form of environmental regulation.

The counterfactual DiD estimate $\beta_{1,Ret}$ is 0.0565 (standard error of 0.0015), compared to the empirical estimate of 0.0440 (standard error of 0.0098). This implies that anticipation biases the Paris Agreement's true effect on firm valuations by 22% (standard error of 0.02). Anticipation produces smaller bias for the true effect on returns than on investment. This is because the return DiD captures changes in two components of the risk premium: the productivity risk premium (which is attenuated by anticipation, as with investment) and the transition risk premium, which is elevated in the *Strict* state because of potential loosening ($\zeta_{Strict,Moderate} > 0$), which inflates Post-period returns in the empirical case relative to the counterfactual, partially offsetting the attenuation.

In comparison, we estimate that the anticipation bias for investment implied by the analytical formula of Hennessy and Strebulaev (2020) is 64%, i.e., more than double the 29%

estimate that our procedure produces.[22] The analytical formula produces a larger bias estimate because it does not account for financing frictions and precautionary savings motives, which partially offset the investment response to anticipated carbon costs and reduce the degree of pre-adjustment. This indicates that our integration procedure provides complementary information for refining causal effect estimates. (The formula does not extend to stock returns, which are not an accumulation variable.) Another benefit of our structural approach is that it can quantify the uncertainty in the bias estimate. Finally, the HS calculation itself requires as inputs the estimated transition intensities which come from our structural estimation, so the formula is not an independent check but a reduced-form projection of the structural model onto the linear-quadratic benchmark it was derived under. Our structural counterfactual remains internally consistent with the full nonlinear model including financing frictions and precautionary savings motives, and it simultaneously corrects for anticipation bias across all moments while preserving the economic mechanisms that generate the differential attenuation across decision variables.

5 Application of the framework in other settings

Our integration procedure can be applied to a wide range of unprecedented but anticipated policy events, including some that are unrelated to environmental regulation. This section describes the steps needed to do so, and proposes two ideas for generalizing our procedure. We focus on applications that require adding only a few steps to standard reduced-form studies.

Consider a reduced-form study of a policy event that was moderately or highly anticipated. Our integration procedure can estimate transition probabilities, given the following inputs: (i) an economic model of the policy setting, with a possibility of state changes; (ii) a numerical algorithm that repeatedly simulates the model and estimates unknown param-

[22] Hennessy and Strebulaev (2020) relate the observed DiD on a firm's accumulation decision (ATE) to its true causal effect (ACE) through the ratio of transient to permanent shadow values of capital: ATE/ACE $= \Delta q/\Delta q_\infty$, where $\Delta q_\infty = (b_4 - b_1)/(r - \hat{\mu}_P + \delta)$ is the permanent shadow value gap if the realized state lasted forever, and $\Delta q = q_4 - q_1$ solves the four-state augmented transition system using the estimated Markov intensities. Since cash flows in our model scale with productivity P_t, the relevant discount rate is $r - \hat{\mu}_P + \delta = 0.2084$, rather than $r + \delta$. Using our estimated transition intensities and tax schedule, the permanent gap is $\Delta q_\infty = -0.1996$ and the transient gap is $\Delta q = -0.0725$, yielding ATE/ACE $= 0.36$ and an implied bias of 64%.

eter values; (iii) a list of moments along with their empirical values in the data; and (iv) a mapping of the realized policy event to one states of the model. Once researchers have obtained transition probability estimates, they can choose from several approaches to adjust reduced-form estimates for anticipation bias, and to provide information about the policy's true causal effects.[23]

For most reduced-form researchers, the economic model and numerical algorithm are the most challenging inputs to provide. Each type of policy event may require its own modeling framework, and it is not straightforward to incorporate state changes into benchmark models used in the literature. Regarding simulation, many economic models—especially dynamic models without closed-form solutions—are computationally intensive to solve. However, key steps of the simulation can be encoded in a generalized script, which estimates a model using researcher-provided data on reduced-form estimates and other empirical moments. This would allow researchers to automate much of the model estimation process.[24] Regarding moments, most are straightforward to measure from widely available data, though prior experience with model estimation can help researchers to identify appropriate moments.[25]

Thus, the key constraint for generalizing our procedure is the economic model. A generalizable model needs to be simple enough for reduced-form researchers to understand which moments to use for estimation. It also needs to be applicable to numerous types of policy events—so that researchers do not need to write a new model for each study—and to incorporate different firm types that correspond to the treatment and control groups of the reduced-form part of the study.

We propose two ideas that may fit these criteria. First, our benchmark model of car-

[23] A simple approach is to verbally evaluate the severity of anticipation bias, and to provide guidance on whether reduced-form estimates constitute an upper or lower bound on the true effect. In some cases, reduced-form estimates can be scaled directly by transition probabilities. Further, causal effects can be credibly estimated using counterfactual analysis, which requires solving a model but is much less computationally intensive than structural estimation.

[24] We are currently developing such a script to potentially disseminate to researchers who want to use our integration procedure. Besides estimating model parameters, a script can also execute comparative static exercises to help researchers identify candidate moments.

[25] Proper identification in model estimation requires that, for each unobservable model parameter, the researcher chooses at least one moment that is highly sensitive to changes in the parameter's value. For example, we verify that our DiD moments are highly sensitive to the transition intensities (see Figure 1). Proper moment selection requires a thorough understanding of the economic model's key frictions and the overall structure of the actual data.

bon taxes is a simple extension of a standard dynamic investment model commonly used in finance (Hayashi, 1982; Abel and Eberly, 1994; Cooper and Haltiwanger, 2006). Carbon taxes are represented by a single parameter that affects firm profits in a linear and additive manner. This parameter could also represent numerous other costs or frictions that firms face. Thus, our model can be applied to several other unprecedented policy events that affect corporate investment, such as the imposition of major tariffs due to trade wars, energy price shocks due to major geopolitical shocks, or major technological shocks that reduce operating costs, such as the release of agentic AI models. Our benchmark model can be used for policy events that do not significantly affect firm financing costs, and it produces closed-form solutions and is computationally easy to estimate.

As one example, Table IA.4 illustrates how our procedure could be used to refine estimates of the effect of tariffs on corporate investment. At the start of the first Trump presidency, firms likely formed beliefs about the possible level of future tariffs, and took these beliefs into account when making investment decisions. In our model from Section 2.2, the carbon tax parameter can be re-defined as a cost that represents tariffs, fixed costs of exporting, or supply chain rigidities. The cost should vary across different trade policy states, and across firms that derive high versus low profits from international trade. Model simulation could show how firms optimally adjust investment, when transitioning from an initial state with no tariffs to one of the states with tariffs. Transition intensities into various tariff regimes could be identified using reduced-form DiD estimates of investment spending around tariff announcements.

Second, our integration procedure could be applied to a political economy model of policy adoption. Researchers could first develop a model of optimal corporate political spending (donations or lobbying), given initial uncertainty about the policy's possible outcomes. Transition probabilities into different policy states could then be obtained by estimating the model, using reduced-form estimates of changes in corporate spending around political events as a key identifying moment. For example, if green firms significantly increase lobbying following the election of a climate-friendly politician, this could indicate high anticipation of a future carbon tax. A key advantage is that political spending is an observable and relevant outcome for many types of policy events, and hence a single model could be applied to estimate transition probabilities in many settings.[26]

6 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 practical guidance for addressing the gap between true causal effects and conventional estimates.

[26]One candidate is the model of corporate lobbying Kang (2016). This model features competition between coalitions with heterogeneous benefits from a policy, which maps naturally to treatment and control groups in a reduced-form setting. The model also features an exogenous probability that a policy will be adopted, which allows for uncertainty that can be resolved after a regime change. Kang (2016) only models two possible policy outcomes (adoption or non-adoption), and thus must be extended to multiple outcomes.

References

Abel, A. B. and J. C. Eberly (1994). A unified model of investment under uncertainty. *American Economic Review 84*(5), 1369–1384.

Albuquerque, R., M. S. Kruttli, Y. Ma, and S. W. Watugala (2026). 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*.

Catherine, S., T. Chaney, Z. Huang, D. Sraer, and D. Thesmar (2022, August). Quantifying reduced-form evidence on collateral constraints. *Journal of Finance 77*(4), 2389–2428.

Cooper, R. W. and J. C. Haltiwanger (2006). On the nature of capital adjustment costs. *Review of Economic Studies 73*(3), 611–633.

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 172*, 104105.

Eisfeldt, A. and G. Schubert (2026). AI and finance. *Annual Review of Financial Economics*, forthcoming.

Eisfeldt, A., G. Schubert, M. B. Zhang, and B. Taska (2026). Generative AI and firm values. *Journal of Finance*, forthcoming.

Erickson, T. and T. M. Whited (2000). Measurement error and the relationship between investment and Q. *Journal of Political Economy 108*(5), 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*(8), 3922–3972.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica 50*(1), 213–224.

Heath, D., M. C. Ringgenberg, M. Samadi, and I. M. Werner (2023). Reusing natural experiments. *Journal of Finance 78*(4), 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.

Hennessy, C. A. and T. M. Whited (2007). How costly is external financing? Evidence from a structural estimation. *Journal of Finance 62*(4), 1705–1745.

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

Kacperczyk, M. T. (2026). What is the carbon premium a premium on? *SSRN Working Paper*.

Kang, K. (2016). Policy influence and private returns from lobbying in the energy sector. *Review of Economic Studies 83*(1), 269–305.

Klausmann, J., P. Krueger, and P. Matos (2025). The green transition: Evidence from corporate green revenues. *SSRN Working Paper*.

McKeon, S. B. (2013). Employee option exercise and equity issuance motives. *SSRN 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. (2026). Carbon pricing versus green finance. *Journal of Finance 81*(2), 561–602.

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 (2026). Climate regulatory risks and corporate bonds. *Journal of Financial Economics*, forthcoming.

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

The Economist (2015). The Paris agreement marks an unprecedented political recognition of the risks of climate change. *The Economist.* December 12, 2015.

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

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

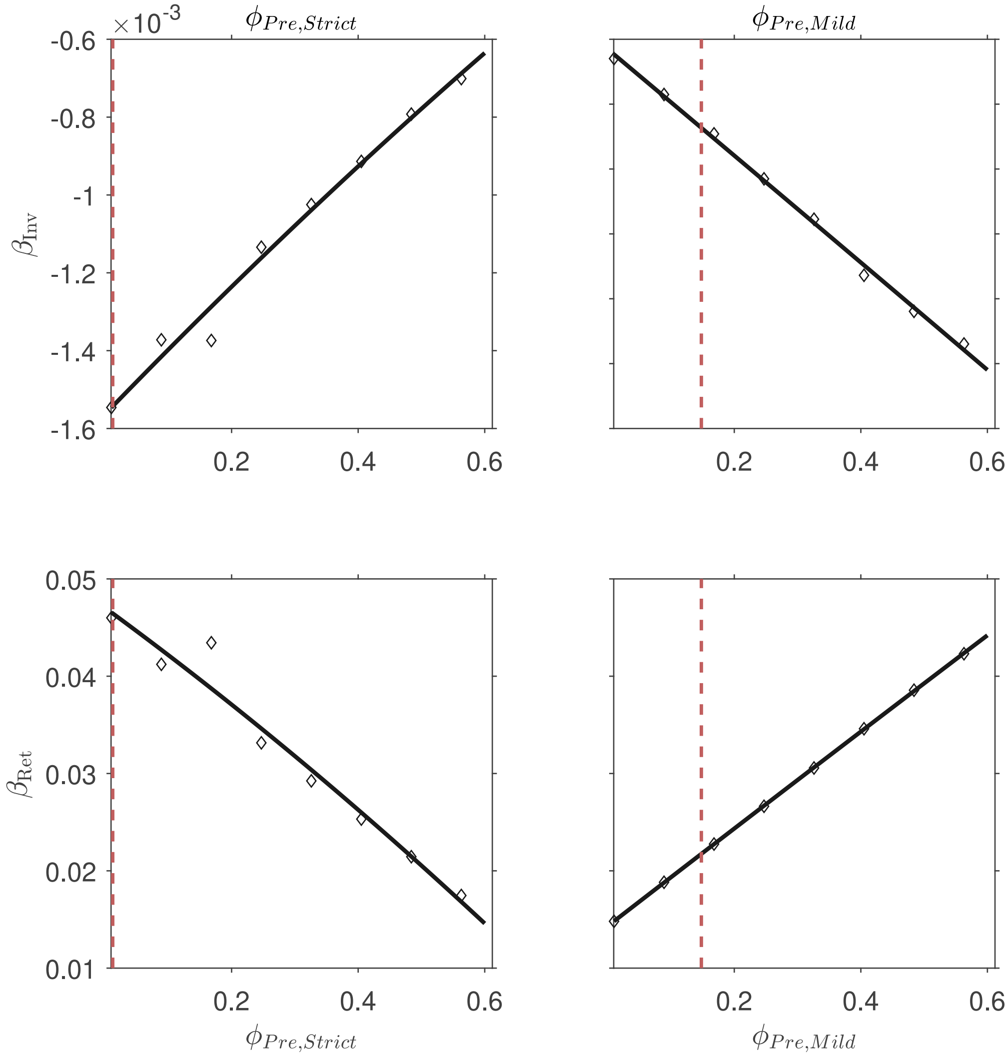

Figure 1: Comparative statics. This figure reports coefficient estimates for $Treated \times Post$, obtained by estimating Eq. (22) using a dataset of simulated model outcomes. For each panel, a set of 50 simulations is run on the full model from Section 2.3, with the transition probability $\phi_{Pre,Mild}$ or $\phi_{Pre,Strict}$ set to the corresponding value on the x-axis. The transition probability $\phi_{Pre,Strict}$ represents the probability of a direct jump from the *Pre* state to the *Strict* state, conditional on leaving the *Pre* state. As this transition probability varies along the x-axis, the total exit intensity out of state *Pre*, as well as the ratio between the intensities from state *Pre* to the *Mild* and *Moderate* states is maintained constant. The transition probability $\phi_{Pre,Mild}$ across these comparative statics is computed similarly. The unknown model parameters are set to the estimated values reported in Table 3. The dashed red line denotes the baseline which is the estimated value of the transition probability.

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 622 U.S. firms in the intersection of Compustat, CRSP, and S&P Trucost, and covers the last three months (quarter Q4) of 2014 through Q4 of 2016 (Q4 of 2015 is omitted). *Stock return* is the unlevered quarterly stock return. We first cumulate the raw monthly return (CRSP data item RET) over each quarter. We then unlever the raw return by dividing it by $(1+(1-\tau)D/E)$, where τ is the firm's marginal tax rate (data item $MTR\_BEFORE\_INTEREST$ from John Graham's tax dataset) and D/E is the total debt-to-equity ratio. Total debt is the sum of short-term (Compustat data item $DLCQ$ and long-term ($DLTTQ$) debt, and equity is the firm's stock market capitalization, measured at the end of each quarter as the price per share ($PRCCQ$) multiplied by shares outstanding ($CSHOQ$). *Investment* is the quarterly change in year-to-date capital expenditures (data item $CAPXY$) scaled by total assets (ATQ). *Size* is the natural logarithm of the stock market capitalization. *Market-to-Book* is stock market capitalization plus total debt, all divided by the sum of total debt and book value of equity ($CEQQ$). *Leverage* is total debt scaled by total assets. *ROE* is net income (NIQ) scaled by stock market capitalization. *PPE* is the natural logarithm of property, plant and equipment ($PPENTQ$). *Beta* is the average market beta over the calendar quarter. *Equity issuance* is the amount of equity issued in each calendar quarter ($STTKQ$) scaled by total assets. This value is truncated from above at the 98th percentile of positive observations, and from below at 2% to avoid stock issuance related to the exercise of employee stock options (in line with evidence in McKeon (2013)). Additional variables are used only in propensity score matching: *Mom* is the monthly raw stock return cumulated over the previous 12 months. *HHI* is the Herfindahl-Hirschman index of concentration of the firm's revenues (Compustat Historical Operating Segment data item $SALES$) across all of its business segments. *Volatility* is average of the monthly stock return volatility (CRSP data item VOL) during the quarter. *Sales growth* is the year-on-year change in quarterly sales (Compustat data item $SALEQ$) scaled by market capitalization. *EPS growth* is the year-on-year change in diluted EPS excluding extraordinary items ($EPSFXQ$) scaled by the price per share. The table reports statistics for *Equity issuance* only for quarters in which issuance occurs. All variables except for *Equity issuance* are winsorized at the 2.5% and 97.5% levels.

Variable	Mean	St. Dev.	25$^{\text{th}}$ Perc.	Median	75$^{\text{th}}$ Perc.	N
Stock return	0.014	0.123	-0.053	0.008	0.077	3,872
Investment	0.014	0.012	0.006	0.011	0.018	5,144
Size	8.647	1.577	7.780	8.651	9.577	5,398
Market-to-Book	2.106	1.355	1.216	1.674	2.492	5,232
Leverage	0.301	0.169	0.189	0.291	0.401	5,258
ROE	0.002	0.040	0.003	0.011	0.017	5,395
PPE	7.607	1.857	6.367	7.633	8.983	5,377
Beta	0.246	0.143	0.136	0.225	0.341	3,966
Equity issuance	0.021	0.044	0.005	0.007	0.013	465
Mom	0.019	0.091	-0.029	0.020	0.065	3,964
HHI	0.251	0.154	0.141	0.224	0.333	3,854
Volatility	0.095	0.059	0.057	0.076	0.111	3,969
Sales growth	-0.021	0.079	-0.028	-0.001	0.010	5,389
EPS growth	-0.003	0.041	-0.008	0.000	0.005	5,402

Table 2: DiD regressions in actual data. This table reports results from DiD regressions (Eq. (22)) estimated using actual data. The sample is 622 U.S. firms in the intersection of Compustat, CRSP, and S&P Trucost, and covers the last three months (quarter Q4) of 2014 through Q4 of 2016 (Q4 of 2015 is omitted). The dependent variable is *Stock return* in columns (1) and (2), and *Investment* in columns (3) through (5). *Post* equals 1 for Q1 through Q4 of 2016, and 0 for Q4 of 2014 through Q3 of 2015. *Treated* equals 1 for firms whose ratio of Scope 1 emissions to revenues (Trucost data item DI_319413 divided by Compustat data item $SALE$) ranks above the 66$^{\text{th}}$ percentile of the distribution in 2014, and 0 for a matched set of firms whose emission intensities are below the 66$^{\text{th}}$ percentile. In columns (4) and (5), we control for lagged values of investment from the previous three quarters. Standard errors are clustered by firm and presented in parentheses. ***, **, * indicate significance levels of 1%, 5%, and 10%, respectively.

Outcome	*Stock return*		*Investment*		
	(1)	(2)	(3)	(4)	(5)
Treated × *Post*	0.0440***	0.0409***	−0.0023***	−0.0015***	−0.0013***
	(0.0098)	(0.0086)	(0.0008)	(0.0004)	(0.0004)
Post	0.0351***	0.0227***	−0.0011***	−0.0007***	−0.0003
	(0.0061)	(0.0065)	(0.0003)	(0.0002)	(0.0002)
Size		−0.1160***			0.0012**
		(0.0223)			(0.0006)
Market-to-Book		−0.0674***			0.0004
		(0.0188)			(0.0003)
Leverage		−0.0776			−0.0106***
		(0.0669)			(0.0030)
ROE		0.3079***			0.0126**
		(0.0864)			(0.0053)
PPE		−0.0775***			−0.0004
		(0.0254)			(0.0010)
Beta		−0.0187			
		(0.0274)			
Lagged investment	No	No	No	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes
N	3,421	2,888	4,563	4,537	4,393
R-squared	0.1599	0.3061	0.7622	0.7798	0.7849

Table 3: Model parameter estimates. This table presents estimates for each unknown model parameter in the set ψ, obtained by estimating the full model from Section 2.3 following our integration procedure as described in Section 4. The model is estimated using SMM. Parameters in Panel A are directly estimated. In Panel B, the parameters for green firms are calculated by linearly interpolating between the estimates for γ_{Pre} and $\gamma_{G,Strict}$, while the parameters for brown firms are calculated by linearly interpolating between the estimates for γ_{Pre} and $\gamma_{B,Strict}$.

Description	**Symbol**	**Estimate**	**Std. Error**
Panel A. Estimated parameters			
(1) Transition intensity from *Pre* to *Mild*	$\zeta_{Pre,Mild}$	0.8049	(0.0010)
(2) Transition intensity from *Pre* to *Moderate*	$\zeta_{Pre,Moderate}$	0.3762	(0.0020)
(3) Transition intensity from *Pre* to *Strict*	$\zeta_{Pre,Strict}$	0.0001	(0.0024)
(4) Transition intensity from *Mild* to *Moderate*	$\zeta_{Mild,Moderate}$	0.0010	(0.0023)
(5) Transition intensity from *Moderate* to *Mild*	$\zeta_{Moderate,Mild}$	1.5139	(0.0006)
(6) Transition intensity from *Moderate* to *Strict*	$\zeta_{Moderate,Strict}$	0.0034	(0.0020)
(7) Transition intensity from *Strict* to *Moderate*	$\zeta_{Strict,Moderate}$	0.3651	(0.0012)
(8) Investment adjustment costs for green firms	θ_G	18.2948	(0.0018)
(9) Investment adjustment costs for brown firms	θ_B	62.2435	(0.0001)
(10) Marginal cost of equity issuance in *Pre* state	γ_{Pre}	0.1871	(0.0002)
(11) Marginal cost of equity issuance for green firms in *Strict* state	$\gamma_{G,Strict}$	0.0353	(0.0001)
(12) Marginal cost of equity issuance for brown firms in *Strict* state	$\gamma_{B,Strict}$	0.4584	(0.0002)
Panel B. Interpolated parameters			
(13) Marginal cost of equity issuance in *Mild* state (green)	$\gamma_{G,Mild}$	0.1365	—
(14) Marginal cost of equity issuance in *Moderate* state (green)	$\gamma_{G,Moderate}$	0.0859	—
(15) Marginal cost of equity issuance in *Mild* state (brown)	$\gamma_{B,Mild}$	0.2776	—
(16) Marginal cost of equity issuance in *Moderate* state (brown)	$\gamma_{B,Moderate}$	0.3680	—
Overidentification test: J-statistic = 115,985,434.11 $\chi^2(6)$ p-value = 0.000			

Table 4: Implied annual transition probability matrix. This table reports the annual transition probabilities across the different regulatory states. This matrix is calculated over discrete time interval dt as $e^{\mathbf{Z}\cdot dt}$, where $\mathbf{Z}$ is the matrix of transition intensities reported in Table 3. Each row gives the probability of transitioning from the current state (denoted by the row label) to each possible state (denoted by column labels) over a one-year horizon.

	Pre	*Mild*	*Moderate*	*Strict*
Pre	0.3069	0.5945	0.0983	0.0003
Mild	0.0000	0.9995	0.0005	0.0000
Moderate	0.0000	0.7788	0.2198	0.0014
Strict	0.0000	0.1550	0.1505	0.6944

Table 5: Simulated and empirical moments. This table lists the 18 moments that the SMM attempts to match during model estimation, reporting their empirical and simulated values. Standard errors are reported in parentheses. Empirical standard errors are computed from the influence function with firm-level clustering. Simulated standard errors are computed via the delta method. The table also reports t-statistics that of the null hypothesis that the empirical moment $\bar{m}_i^{\text{emp}}$ equals the simulated moment $\bar{m}_i^{\text{sim}}$. The t-statistics are computed as $(\bar{m}_i^{\text{emp}} - \bar{m}_i^{\text{sim}})/\widehat{\text{SE}}(\bar{m}_i^{\text{emp}})$. Panels A and B report moment values multiplied by 100, while Panel C reports values multiplied by 10,000. In Panel A, $\beta_{1,Inv}$ and $\beta_{1,Ret}$ are the coefficient estimates $Treated \times Post$ obtained from regression based on columns (4) and (1) of Table 2.

Description	Empirical	Simulated	t-stat
Panel A. Difference-in-difference moments			
(1) $\beta_{1.Inv}$	−0.152	−0.134	−0.42
	(0.043)	(0.010)	
(2) $\beta_{1.Ret}$	4.399	4.709	−0.32
	(0.980)	(0.310)	
Panel B. Equity issuance moments			
(3) Issuance frequency green firms before	1.184	1.245	−0.17
	(0.367)	(0.290)	
(4) Issuance frequency green firms after	0.941	4.545	−9.76
	(0.369)	(0.350)	
(5) Issuance frequency brown firms before	2.224	2.002	0.42
	(0.526)	(0.750)	
(6) Issuance frequency brown firms after	3.809	3.193	0.79
	(0.777)	(0.440)	
(7) Issuance amount green firms before	3.906	5.454	−1.33
	(1.161)	(0.820)	
(8) Issuance amount green firms after	4.485	5.428	−0.15
	(6.315)	(0.150)	
(9) Issuance amount brown firms before	4.617	5.551	−1.07
	(0.876)	(0.000)	
(10) Issuance amount brown firms after	4.724	4.336	0.37
	(1.052)	(0.350)	
Panel C. Second moments			
(11) Variance of issuance amounts green firms	1.972	3.275	−0.08
	(15.588)	(0.500)	
(12) Variance of issuance amounts brown firms	1.628	2.656	−0.38
	(2.729)	(0.500)	
(13) Variance of investment green firms	0.209	0.003	7.62
	(0.027)	(0.020)	
(14) Variance of investment brown firms	0.551	0.004	9.46
	(0.058)	(0.020)	
(15) Covariance of cash and investment green before	−0.047	0.019	−1.07
	(0.062)	(0.020)	
(16) Covariance of cash and investment green after	−0.095	0.021	−1.76
	(0.066)	(0.010)	
(17) Covariance of cash and investment brown before	−0.342	−0.020	−3.12
	(0.103)	(0.010)	
(18) Covariance of cash and investment brown after	−0.287	−0.040	−2.49
	(0.099)	(0.030)	

Table 6: Anticipation bias in DiD moments. This table reports the coefficient estimates $Treated \times Post$, obtained from using various panel datasets to run regressions based on columns (4) and (1) of Table 2. The column *Counterfactual* reports estimates from a counterfactual dataset obtained from simulated the model with all transition intensities in ψ set to zero (no anticipation), and other parameters set to the values in Table 3. The column *Empirical* reports estimates from the actual data. Standard errors are in parentheses. We compute simulation standard errors for *Counterfactual* and *Estimated* (based on 5 independent replications of 50 simulations each) and firm-clustered standard errors for *Empirical*. Δ is the difference between *Counterfactual* and *Empirical*, with standard errors in parentheses calculated as $\sqrt{\mathrm{SE}^2_{Counterfactual} + \mathrm{SE}^2_{Empirical}}$. *Bias* is calculated as *Bias* = 1− *Empirical*/*Counterfactual* and thereby measures the fraction of the causal DiD coefficient attenuated by anticipation, with standard errors based on the delta method in parentheses. Asterisks on *Estimated* indicate whether it differs statistically from *Empirical*: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Moment	*Counterfactual*	*Empirical*	Δ	*Bias*
$\beta_{1,Inv}$	-0.0021	-0.0015	-0.0006	0.29
	(0.0001)	(0.0004)	(0.0004)	(0.03)
$\beta_{1,Ret}$	0.0565	0.0440	0.0125	0.22
	(0.0015)	(0.0098)	(0.0099)	(0.02)

Internet Appendix

for

Addressing Anticipation Effects in Finance

This internet appendix provides a calibration of the model with the goal to show how measured treatment effects can deviate from causal effects, a general overview of the SMM approach, and additional tables.

IA.1 Calibration

IA.1.1 Setting

In this internet appendix, we provide a calibration of 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 $\gamma_{B,1} = \gamma_{G,1} = 6\%$ 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: $\gamma_{B,2} = \gamma_{G,2} = 6\%$. 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 $\gamma_{G,3} = 6\%$; the fixed cost of equity issuance for brown firms is $\phi_{B,3} = 2\%$ and the marginal cost is $\gamma_{B,3} = 6\%$. 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

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

transition intensities to these states to: $\zeta_{(1,2)} = \zeta_{(1,3)} = 0.3$. To simplify the analysis, we 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.

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.1.2 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 1 for brown firms, and 0 otherwise, and $Post$ equals 1 when the state has switched out of State 1, and 0 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.1.1, 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.1.3 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 IA.6 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.

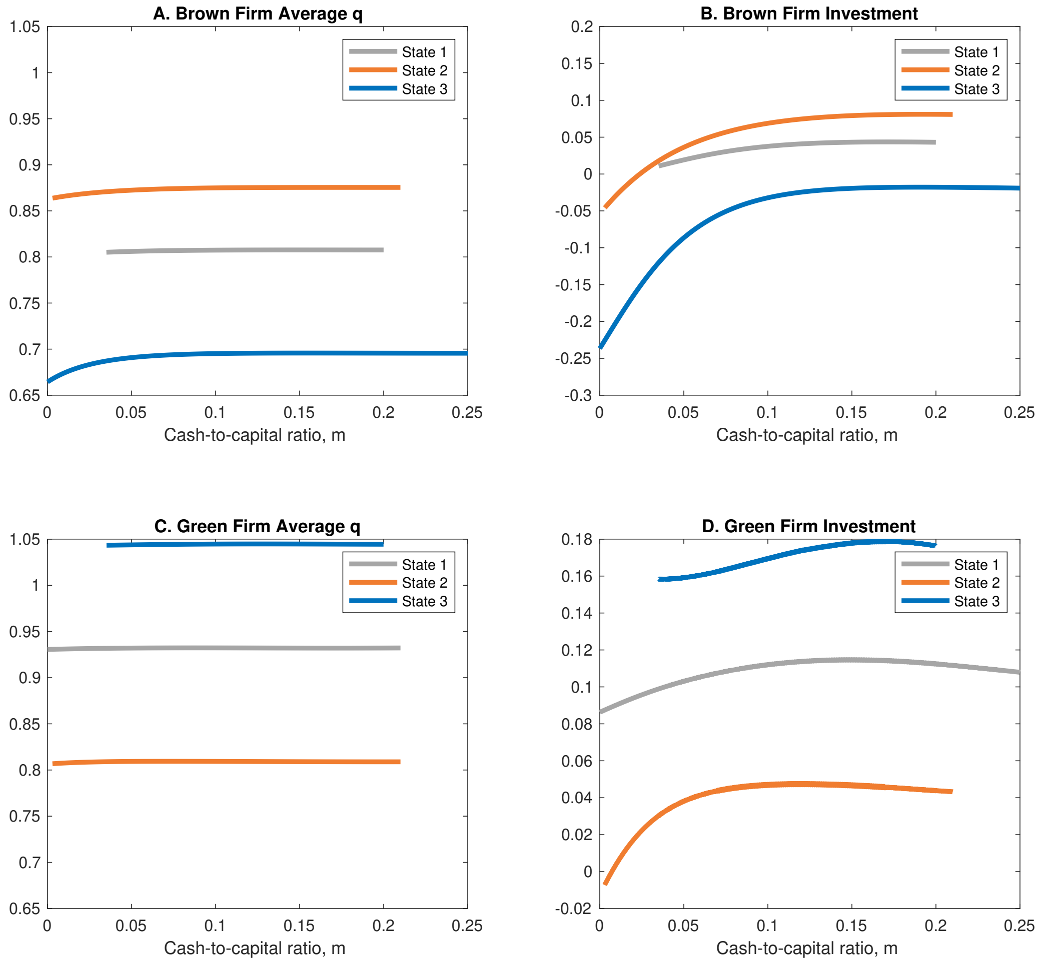

Figure IA.1: Average q and investment. The figures show how average q and investment vary with the cash-to-capital ratio for representative green and brown firms, given the climate regulation state. State 1 denotes the state where there are no carbon taxes. The costs of capital in State 1 are identical 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.

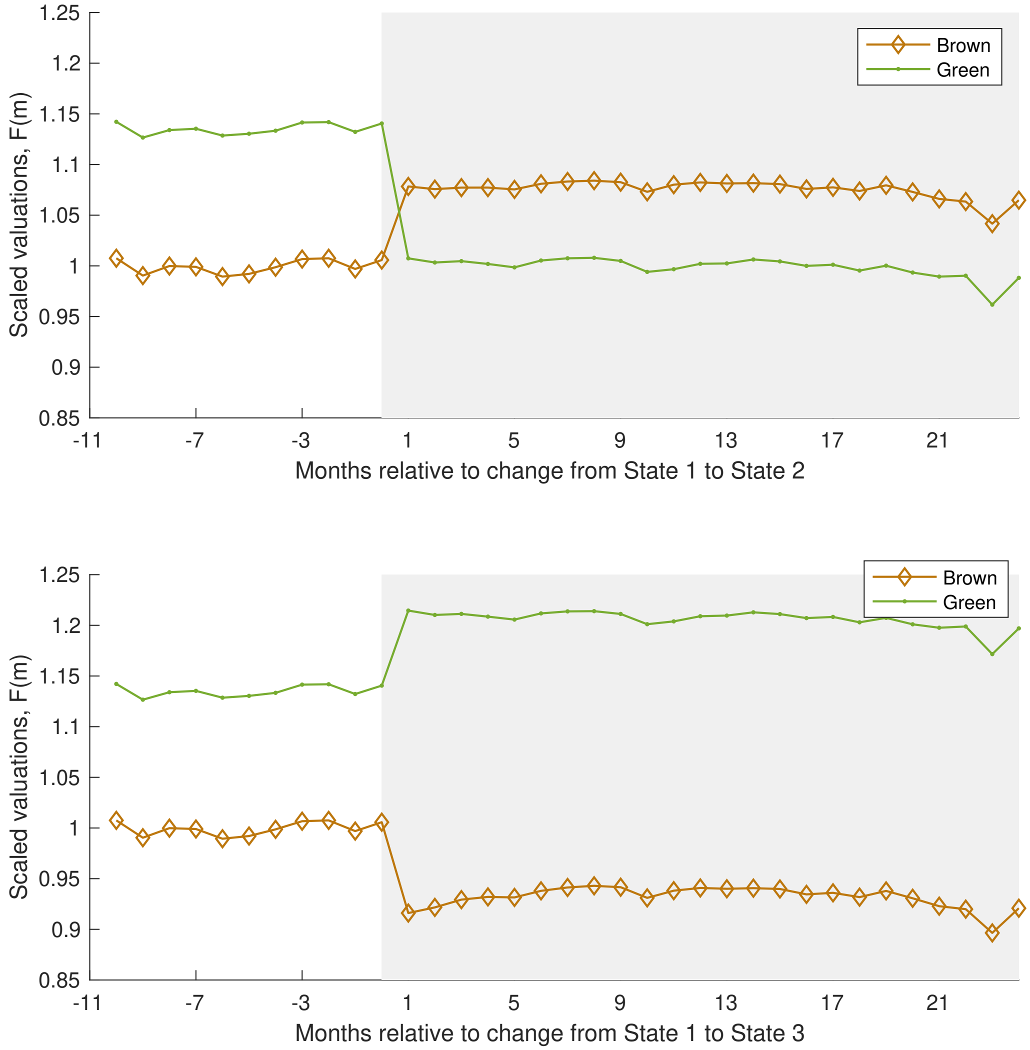

Figure IA.2: Green vs. brown firms. This figure shows how the (cross-sectional) average scaled valuations of green and brown firms change in the months around the change out of State 1. The top panel shows these changes when State 2 is realized; the bottom panel shows the changes in value that occur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

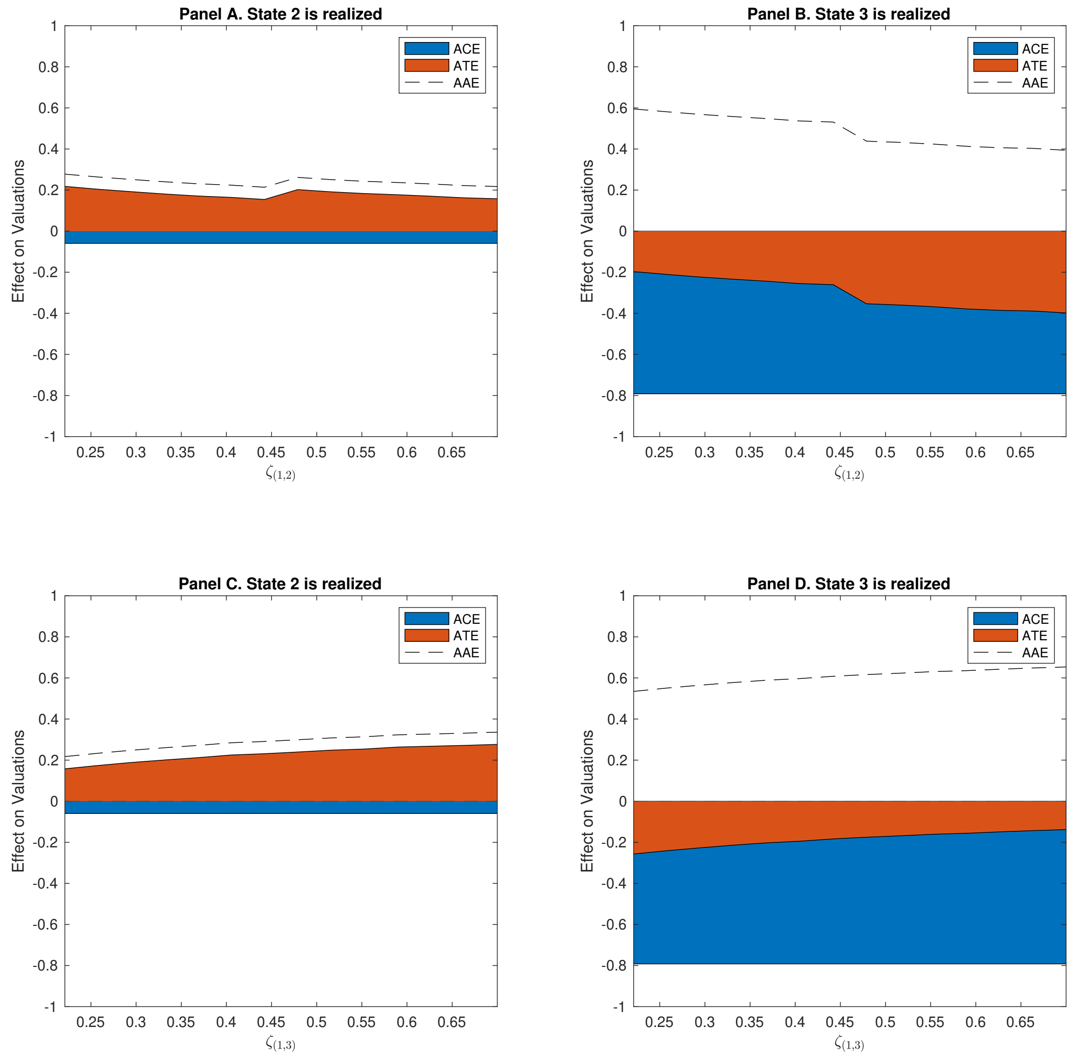

Figure IA.3: Average treatment, causal, and anticipation effects. The figures show how treatment, causal and anticipation effects vary with the transition intensities out of State 1. The top two panels show the variation with respect to the level of the carbon tax in State 2. The first panel shows the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

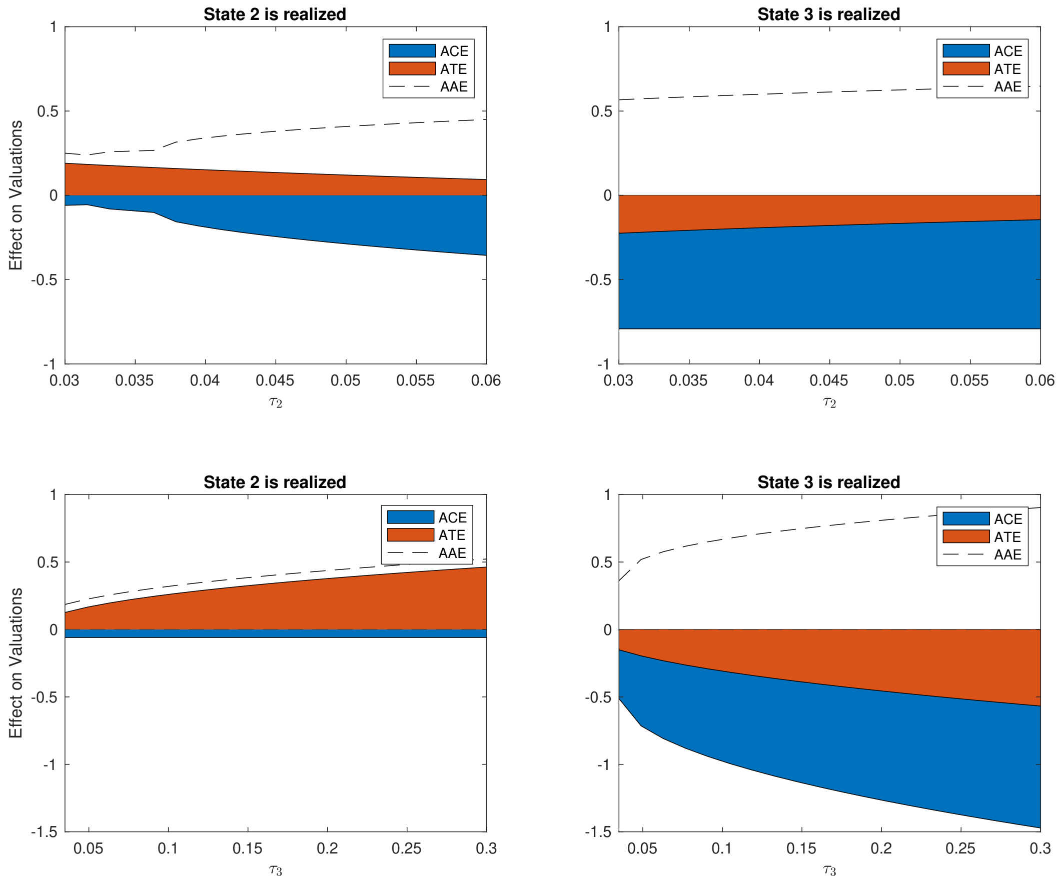

Figure IA.4: Average treatment, causal, and anticipation effects. The figures show how treatment, causal and anticipation effects vary with the level of carbon taxes. The top two panels show the variation with respect to the level of the carbon tax in State 2. The first panel shows the effects when State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

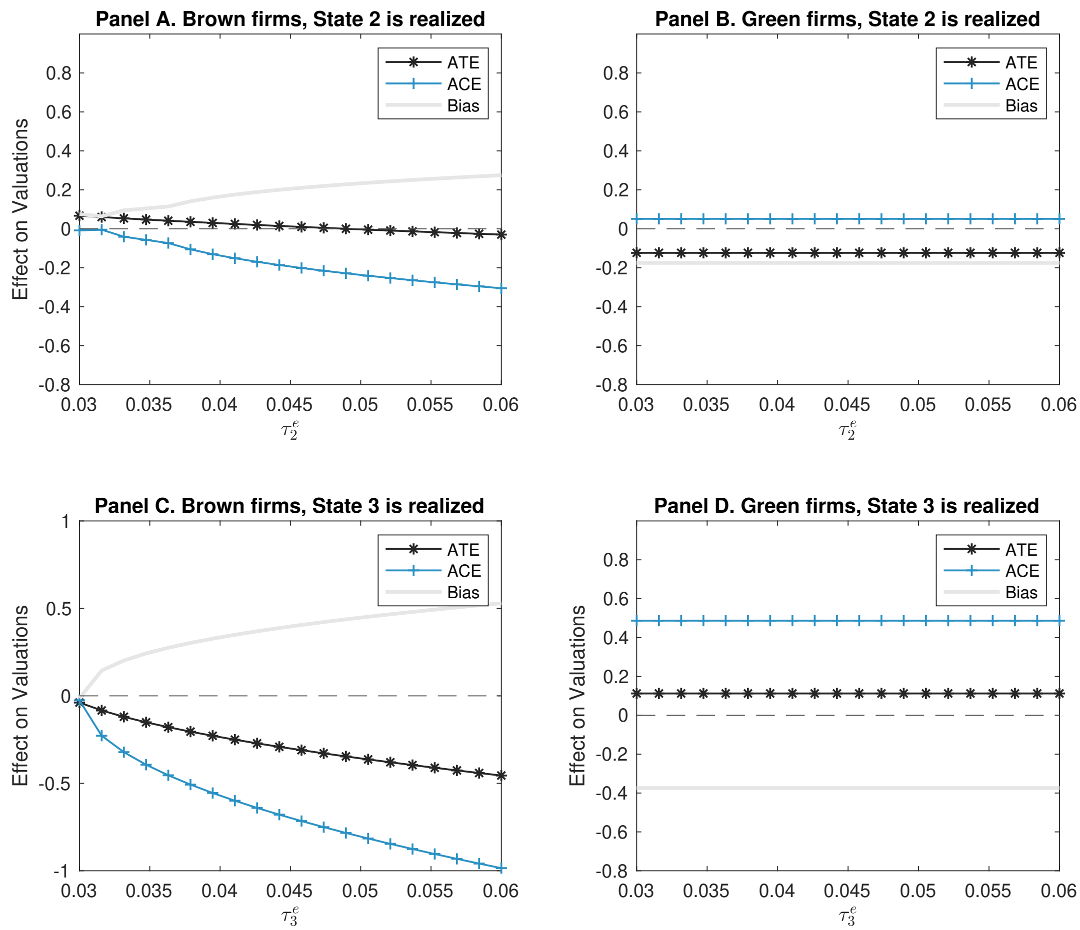

Figure IA.5: Comparative statics: Carbon tax rates. The figures show how the average treatment and causal effects, as well as the bias in treatment effects induced by anticipation vary with the level of carbon taxes. Panels A and B show these comparative statics with respect to the carbon tax rate in 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.

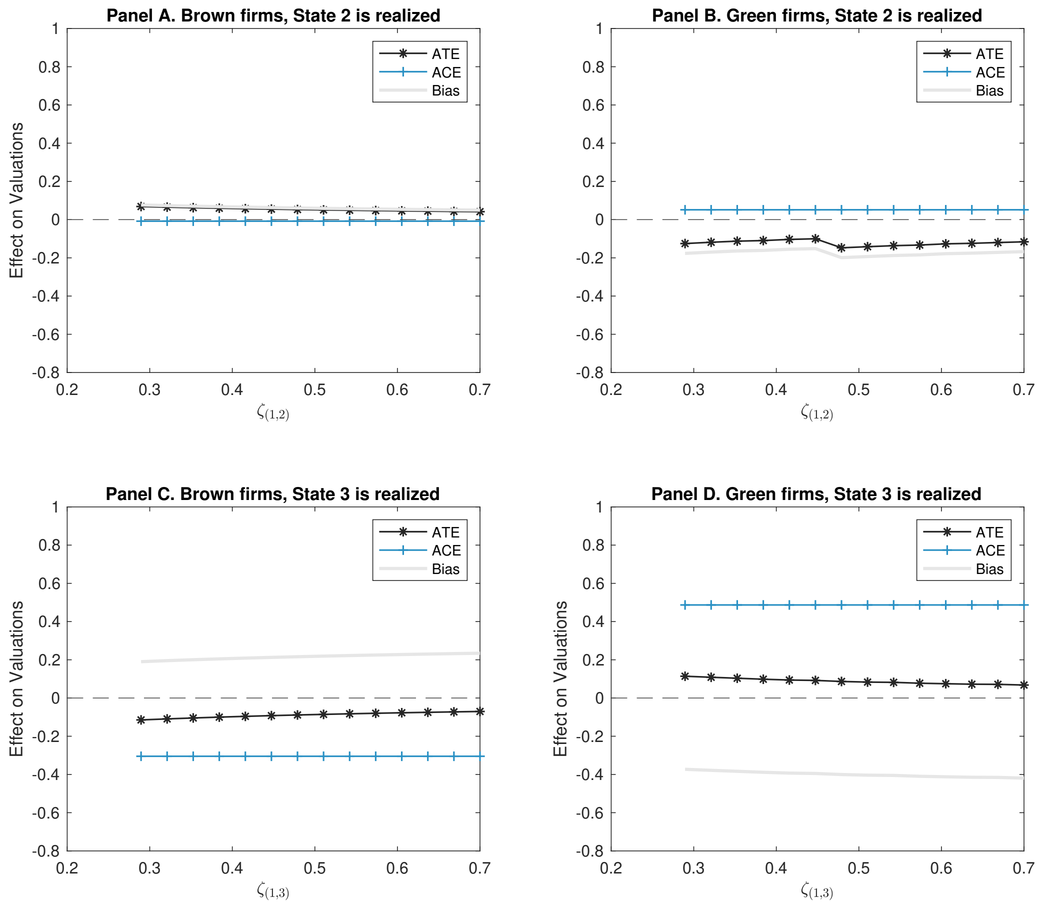

Figure IA.6: Comparative statics: Transition intensities out of State 1. The figures show how the average treatment and causal effects, as well as the bias in treatment effects induced by anticipation vary with the level transition intensities out of State 1. Panels A and B show these comparative statics 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.2 General overview of SMM

In this internet appendix, we provide a general overview of the SMM approach. It is targeted primarily to reduced-form researchers that might be less familiar with the approach. 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 . 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.

IA.3 Additional tables

Table IA.3: DiD regressions in actual data—Alternative treatment indicator This table reports results from DiD regressions (Eq. (22)) estimated using actual data. The sample is 462 U.S. firms in the intersection of Compustat, CRSP, and S&P Trucost, and covers the period from Q4 of 2014 through Q4 of 2016. The dependent variable is *Stock return* in columns (1) and (2), and *Investment* in columns (3) through (5). *Post* equals 1 for Q1 through Q4 of 2016, and 0 for Q1 through Q3 of 2015 (Q4 of 2015 is omitted). *Treated2* equals 1 for firms whose ratio of Scope 1 emissions to revenues in 2014 ranks above the 75^{th} percentile of the sample distribution, and 0 for a matched set of firms whose emission intensities are below the 75^{th} percentile. For *Investment*, lagged investment controls include the first three quarterly lags of *Investment*. Additional controls include lagged values of $Size$, *Market-to-Book*, $Leverage$, $Cash$, ROE, and PPE. For *Stock return*, controls include lagged values of $Size$, *Market-to-Book*, $Beta$, $Leverage$, $Cash$, ROE, and PPE. All specifications include firm fixed effects. Standard errors are presented in parentheses. ***, **, * indicate significance levels of 1%, 5%, and 10%, respectively.

Outcome	*Stock return*		*Investment*		
	(1)	(2)	(3)	(4)	(5)
Treated2 × *Post*	0.0398**	0.0373**	−0.0028**	−0.0017***	−0.0014**
	(0.0145)	(0.0126)	(0.0008)	(0.0004)	(0.0004)
Lagged investment	No	No	No	Yes	Yes
Additional controls	No	Yes	No	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes
N	3,105	2,611	3,378	3,352	3,244
R-squared	0.1617	0.3069	0.7344	0.7556	0.7628

Table IA.4: Extension of methodology to other settings. This table illustrates step-by-step how our approach can be used to examine the role of anticipation effects in a trade war setting.

Step	**Our Paper: Paris Agreement**	**Alternative Setting: Trade War**
1. Define policy event	Paris Agreement is treated as a major anticipated climate-policy shock.	A trade-war event could be the announcement or implementation of new tariffs, export controls, sanctions, or retaliatory measures.
2. Define treated and control groups	Treated firms are high-emissions "brown" firms; control firms are lower-emissions "green" firms.	Treated firms are firms highly exposed to tariff-targeted countries or industries; controls are firms with low exposure.
3. Estimate reduced form	Estimate DiD effects around Paris Agreement for outcomes such as investment, cash holdings, and stock returns.	Estimate DiD effects around tariff announcements for imports, exports, investment, inventories, markups, employment, or stock returns.
4. Build structural model	Model embeds carbon taxes, investment adjustment costs, costly external finance, and cash-saving motives.	Model could include tariffs, fixed export/import costs, supply-chain adjustment costs, inventory frictions, and financing constraints.
5. Add anticipation	Firms assign probabilities to future climate-policy states: no regulation, mild, moderate, or strict regulation.	Firms assign probabilities to future trade-policy states: no tariff, mild tariff, broad tariff, retaliation, or escalation.
6. Simulate panel data	Model simulates green and brown firms before and after possible climate-policy transitions.	Model simulate high- and low-trade-exposure firms before and after possible tariff or retaliation states.
7. Re-estimate reduced form in simulated data	Same DiD specification used in the real Paris Agreement data is run on the simulated model-generated data.	Same empirical DiD or event-study design would be run on simulated trade-war data.
8. Match simulated and empirical moments	SMM chooses parameters so that simulated DiD coefficients and other moments match empirical estimates.	SMM chooses tariff-belief probabilities, adjustment costs, and trade frictions so that simulated estimates match observed responses.
9. Recover parameters and beliefs	Recover firms' implied beliefs about the probability of climate regulation before Paris.	Recover firms' implied beliefs about tariff escalation or retaliation before the trade-war announcement.
10. No-anticipation counterfactual	Model is re-solved with pre-event transition probabilities set to zero.	Model is re-solved assuming firms did not anticipate tariffs or retaliation.
11. Quantify anticipation bias	Compare observed reduced-form effects with the no-anticipation causal effect and show that reduced-form estimates understate the Paris effect.	Quantify whether tariff studies understate or overstate true effects because firms adjusted sourcing, inventories, prices, or investment before the announcement.
12. Counterfactual policy analysis	Model can compare alternative climate-policy paths or belief distributions.	Model can compare alternative tariff designs, gradual versus sudden implementation, exemptions, retaliation probabilities, or credible pre-announcement signals.