Time Inconsistency and Reputation in Climate Policy*

Inès Mourelon[†] & Andrea Titton[‡]

This version: June 2026

Abstract

This paper studies how reputation affects climate policy when governments cannot fully commit to future carbon prices. We develop a continuous-time reputation game in which firms choose abatement based on the climate policy they expect to face, while a non-committed government re-optimizes policy sequentially. Firms do not observe the government's type directly and update their beliefs from noisy public signals. Reputation therefore becomes an endogenous state variable that affects both firms' abatement decisions and the government's incentive to deviate from the committed policy path. We characterize the Markov sequential equilibrium as a function of cumulative emissions and policy reputation. A government with a stronger reputation can set a lower current climate policy while keeping firms' expected carbon price and abatement relatively high. This policy flexibility is constrained by learning, as deviations from the committed path reveal information about the government's type and erode its reputation over time. Noisier policy signals weaken this constraint and allow the government to deviate further. The value of reputation is greatest when reputation is already high and cumulative emissions remain relatively low.

JEL Classifications: Q58; D83; C73

Keywords: Climate policy; Reputation; Time inconsistency; Bayesian learning; Continuous time.

---

*The authors gratefully acknowledge the financial support of the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (Grant agreement No 101170509 — BELIEFS)

[†]University of Bologna. Email: ines.mourelon@unibo.it

[‡]University of Bologna. Email: andrea.titton2@unibo.it

1 Introduction

One objective of carbon pricing is to shape investment decisions by making firms internalize the future cost of their emissions today. In the European Union, this logic is embodied in the EU Emissions Trading System (EU ETS), launched in 2005 and which operates through a declining emissions cap and a market for tradable allowances. Yet recent political debates suggest that even a well-established carbon pricing system may be subject to uncertainty about its future path. In early 2026, Italy called for the suspension of the EU carbon market, citing concerns about industrial competitiveness and energy costs, while other European policymakers defended the scheme and its stabilisation mechanisms. More broadly, debates surrounding the extension of emissions trading to buildings and road transport under the new ETS2 highlight that the future path of carbon pricing remains politically contested.[1]

In this context, how do firms choose abatement when they are uncertain about the future carbon price, and about the persistence of the policy itself? This question is especially important because carbon-policy decisions have both immediate transition costs and long-lasting climate consequences. Policymakers may be tempted to deviate from previously announced policy paths when the economic or political costs of maintaining them become high, creating a time-inconsistency problem (Kydland & Prescott, 1977). Anticipating this possibility, firms base their abatement decisions not only on the current climate policy, but also on their beliefs about the government's commitment to future policy. Recent movements in EU carbon prices have been tied in part to political signals about potential reform, suggesting that markets update expectations about the future policy environment in response to political announcements.[2]

This paper studies climate policy through the lens of time inconsistency, reputation, and learning. We develop a dynamic model in which firms make costly abatement decisions that depend on the climate policy they expect to face. These decisions are forward-looking and depend both on the current level of carbon pricing and on firms' beliefs about whether policymakers will maintain the policy in the future. The policymaker may be either committed, in which case it implements a previously announced contingent policy plan, or non-committed, in which case it re-optimizes climate policy sequentially. Firms do not observe the policymaker's type directly. Instead, they receive noisy public signals about future climate policy and update their beliefs over time. We interpret these beliefs as the reputation of climate policy. Signals suggesting a departure from the committed policy path weaken reputation and therefore affect firms' expected climate policy, current abatement, and emissions. Be-

[1] See Euronews, Feb. 2026 and Euronews, March 2026 (last accessed, April 15, 2026).

[2] See Reuters, Feb. 2026 (last accessed, April 15, 2026).

cause emissions accumulate over time, temporary credibility losses can have persistent effects on the climate state and on social welfare.

This paper shows that reputation affects climate policy in the absence of commitment. A non-committed government with a strong reputation can set a lower current climate policy while relying on firms' belief that the committed policy will prevail. Because firms place greater weight on the committed tax, the expected carbon price remains relatively high, sustaining abatement despite a lower tax actually implemented. By contrast, a government with little reputation has less scope to exploit firms' beliefs and must choose a higher current tax, closer to the committed benchmark, to preserve abatement incentives. This strategic use of reputation is constrained by learning: larger deviations from the committed policy make public signals more informative about the government's type, decreasing its reputation over time. The strength of this effect depends on signal precision. When signals are more informative, deviations carry a larger reputational cost; when signals are noisier, the non-committed government can depart further from the committed path without generating an equally strong revision in beliefs. Finally, differences in reputation affect equilibrium taxes and abatement and can therefore carry over to the subsequent path of cumulative emissions.

Our paper's main contribution is to characterize and quantify how reputation affects climate policy through two complementary channels. First, reputation shapes firms' expectations of future carbon pricing and therefore their current abatement decisions. Second, the government internalizes the effect of its current policy on future beliefs, so that reputation acts as an endogenous cost of discretion. By combining this reputational state variable with cumulative emissions, the model shows how temporary policy deviations can alter the long-run transition path.

We add to work on the time inconsistency of climate policy by showing that the effects of discretionary deviations depend on how firms learn about future policy. Climate policy has long been studied as a commitment problem. Following the classic time-inconsistency problem exposed by Kydland and Prescott (1977), policies that are optimal ex ante may no longer be optimal ex post once private agents have adjusted their behavior. This concern is especially relevant in climate policy, where regulation is announced over long horizons but future policymakers may have incentives to revise previously announced paths (Brunner et al., 2012; Helm et al., 2003). Credibility matters because low-carbon investments are often costly and irreversible. Previous work has highlighted that when firms incur sunk costs before policy is fully implemented, governments may be tempted to deviate ex post, creating hold-up problems (Gersbach & Glazer, 1999) and weakening incentives to invest in abatement or cleaner technologies (Acemoglu & Rafey, 2023; Requate & Unold, 2001). These credibility failures can delay the transition and increase the risk of stranded assets (Kalkuhl et al., 2020). Our framework instead emphasizes how beliefs about policy commitment affect a continuous flow of abatement decisions and how the resulting emissions accumulate over time.

A related literature studies reputation as a partial substitute for commitment. Building on the broader theory of dynamic games with incomplete information, where reputation is understood as a belief about an agent's type (Kreps & Wilson, 1982), this strand shows that discretionary authorities may behave more consistently when current actions shape beliefs about future behavior (Barro & Gordon, 1983). While this logic has begun to be applied to climate policy (Chiappinelli & May, 2022), existing work says less about how firms learn about future policy when commitment is not directly observed. This issue becomes central when observed policy signals are noisy or when policymakers have imperfect control, so that credibility evolves gradually through Bayesian learning rather than being updated instantaneously (Canzoneri, 1985; Cukierman & Meltzer, 1986; Faust & Svensson, 2001; King & Lu, 2022; Lu et al., 2016). Our paper builds on these insights by studying climate policy as both a commitment and a reputational problem. Firms make abatement decisions while updating their beliefs about future carbon pricing from noisy signals. Our contribution is to make credibility endogenous: reputation is learned from policy signals and, in turn, shapes firms' investment decisions.

The remainder of the paper is organized as follows. Section 2 introduces the model and Section 3 derives the equilibrium and presents the optimal policies. Section 4 reports the main results. Section 5 concludes.

2 Model

This section introduces a partial equilibrium model of the climate transition costs in which we illustrate the issue of time inconsistency.

2.1 Firms

Time is continuous $t \geq 0$. There is a continuum of polluting firms $i \in [0, 1]$, endowed with a production technology and emitting e_0 GtCO_2e/year in absence of abatement efforts. Firms can reduce their emissions $e_{i,t}$, by abating $a_{i,t}$ GtCO_2e/year, at a quadratic cost $c(a_{i,t}) := \nu a_{i,t}^2/2$ as a fraction of output y_0 t\$/year. Finally, all firms pay a carbon tax τ_t t\$/GtCO_2e on their emissions $e_{i,t}$. The firms' total costs are given by

$$k(a_{i,t}, \tau_t) := (e_0 - a_{i,t})\tau_t + y_0 c(a_{i,t}). \tag{1}$$

Let $e_t := \int_0^1 e_{i,t}\,\mathrm{d}i$ and $a_t := \int_0^1 a_{i,t}\,\mathrm{d}i$ be the aggregate emissions and abatement efforts respectively. As all firms are endowed with the same production technology, the total level of emissions follows $e_t = e_0 - a_t$.

2.2 Welfare and Government

This section introduces transition costs and climate damages. These two cost channels are internalised by the government but not by firms, which are atomistic.

Transition costs are driven by stranded assets. We model stranded assets as a reduce form costs imposed by climate policy, which depend on the abatement efforts of the economy. We consider two possible channels for such costs. First, higher climate policy τ_t lower the market value of existing fossil fuel capital which, if used as collateral by firms, can generate welfare costs via financial frictions (Van Der Ploeg & Rezai, 2020). Second, a higher climate policy forces a quick retirement of polluting capital, thereby generating large adjustment costs bore by society (Baldwin et al., 2020; Helm et al., 2003; Kalkuhl et al., 2020). The former marginal costs are decreasing in the fraction of abated emissions and are modeled as

$$R(a_t) \coloneqq \delta_r e_0 \left(e_0 - a_t\right). \tag{2}$$

The latter marginal costs are increasing in the fraction of abated emissions

$$A(a_t) \coloneqq \delta_a \left(a_t - \underline{a}\right)_+^2, \tag{3}$$

where $\underline{a}$ is a frictionless level of baseline abatement.[3] Together, the total transition losses, in t\$/year, are given by

$$l(\tau_t, a_t) \coloneqq \frac{A(a_t) + R(a_t)}{2}\tau_t^2 \tag{4}$$

Cumulative emissions

$$m_t \coloneqq \int_0^t \int_0^1 e_{i,s}\,\mathrm{d}i\,\mathrm{d}s\,, \tag{5}$$

increase temperature deviations from its pre-industrial level and, as a consequence, create climate damages. Following Dietz and Venmans (2019), we summarise the economic consequences of climate change assuming that cumulative emissions m_t, generate warming ζm_t and, in turn, damages

$$d(m_t) \coloneqq 1 - e^{-\frac{\gamma}{2}(\zeta m_t)^2} \tag{6}$$

[3]Here, $(\cdot)_+ \coloneqq \max\{\cdot, 0\}$.

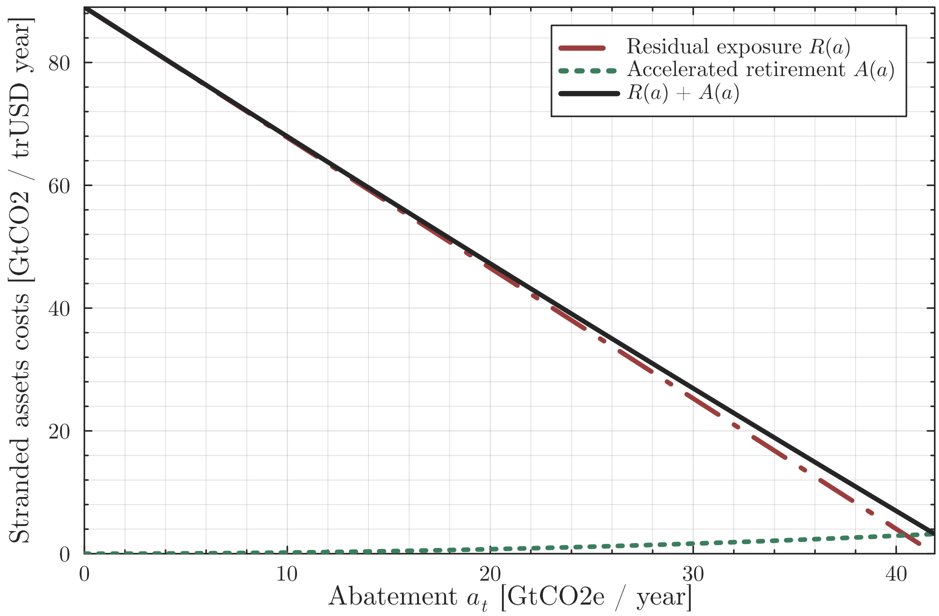

Figure 1: The two channels of marginal transition losses $R(a)$ (dashed) and $A(a)$ (dotted), at different abatement levels a_t. The total marginal transition losses $R(a) + A(a)$ are displayed as a solid line.

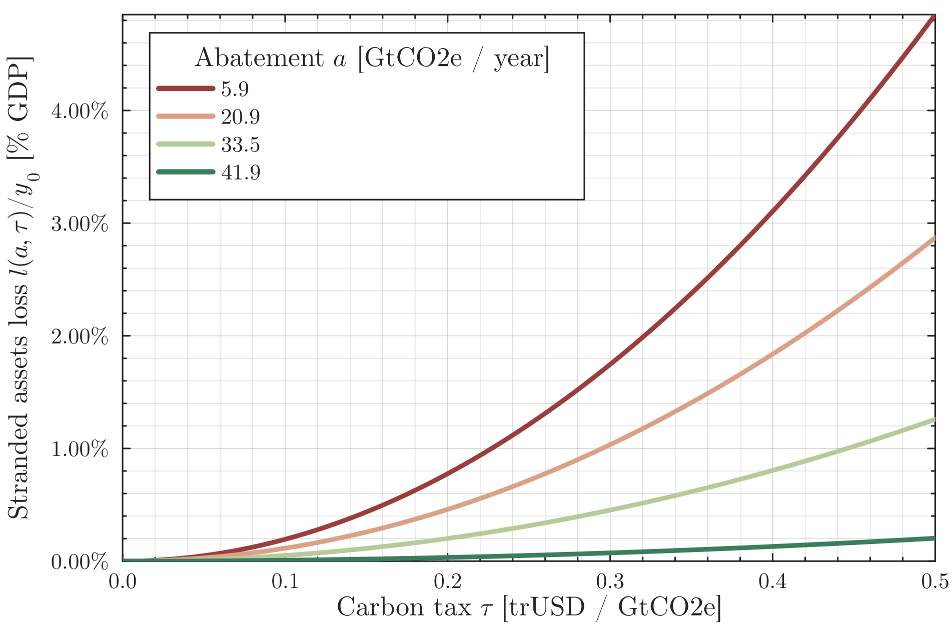

Figure 2: Total transition losses $l(a_t, \tau_t)$ as a function of different carbon taxation levels, for increasing levels of abatement a_t.

as a fraction of output y_0.

The government sets carbon taxes τ_t to minimise the net present value of the welfare costs of the green transition

$$\int_0^\infty re^{-rt} w(\tau_t, a_t)\,\mathrm{d}t \tag{7}$$

where instantaneous welfare costs are given by

$$w(\tau_t, a_t) \coloneqq y_0\,(d(m_t) + c\,(a_t)) + l(a_t, \tau_t). \tag{8}$$

2.3 Policy Signals and Committent

We model time inconsistency by extending the framework introduced in Faingold and Sannikov (2011), to an environment with an additional state variable: cumulative emissions m_t. We assume that the government can be of two types θ, either committed $\theta = \mathrm{c}$ or non-committed $\theta = \mathrm{nc}$. A committed government c chooses, at time $t = 0$, a tax schedule $\tau^c_t = \tau^c(m_t)$ for all $t \geq 0$. A non-committed government nc chooses at each moment in time t, a carbon tax τ_t. Firms do not know the government's type nor the future governments' policy. Firm, receive a noisy signal s_t of carbon taxes, which satisfies

$$\mathrm{d}s_t = -\epsilon(\tau^c - \tau_t)\,\mathrm{d}t + \sigma\,\mathrm{d}b_t\,, \tag{9}$$

where b_t is a Brownian motion. We assume the signal s_t to be a policy factor that firms can extract from the price of futures contracts on emission permits in the secondary market.

As in (Faingold & Sannikov, 2011), we characterise the equilibrium of play between firms and the government as a sequential equilibrium in a reputation game. Firms begin the game with a common prior belief ϕ_0 about the government commitment $\theta = \mathrm{c}$. At each moment in time t, companies make abatement decisions a_t, based on their belief ϕ_t. This belief is updated in turn using the Bayes rule, by observing the signal s_t. As the signal s_t is common and the prior belief is the same for all firms, the evolution of belief ϕ_t is the same for all firms. Bayes' rule implies that the beliefs ϕ_t satisfy the equation

$$\mathrm{d}\phi_t = -\phi_t(1-\phi_t)\frac{\epsilon(\tau^c_t - \tau_t)}{\sigma^2}\left(\mathrm{d}s_t + (1-\phi_t)\epsilon(\tau^c_t - \tau_t)\,\mathrm{d}t\right). \tag{10}$$

The belief dynamics (10) can be further simplified under the non-committed type to

$$\frac{\mathrm{d}\phi_t}{\phi_t(1-\phi_t)} = -\phi_t \xi(\tau_t)^2 \,\mathrm{d}t + \xi(\tau_t)\,\mathrm{d}Z_t \tag{11}$$

$$\xi(\tau) \coloneqq \frac{\epsilon}{\sigma}\big(\tau^c - \tau\big). \tag{12}$$

Importantly, as long as $\tau_t < \tau^c$, the belief process is a sub-martingale, that is, $\mathbb{E}_t \phi_s \leq \phi_t$ for all $s \geq t$. This implies that the firms will eventually recognise a non-committed government who chooses a tax τ_t lower than the committed one τ_t^c.

3 Equilibrium & Optimal Policies

The game presented above satisfies Condition 1 and 2 in (Faingold & Sannikov, 2011, p.797) for a unique Markovian sequential equilibrium. In this section, we present and compute the optimal strategies of the three players: firms, the committed c government, and the non-committed nc government.

3.1 Firms

Given beliefs ϕ and a non-committed tax τ, the firms' first order condition on their costs (1), implies that the optimal abatement best response is given by

$$a(\phi, \tau) \coloneqq \min\left\{ \frac{\phi \tau^c + (1-\phi)\tau}{y_0 \nu}, e_0 \right\}. \tag{13}$$

As its costs are quadratic, the firm computes the expected climate policy given its beliefs $\phi\tau^c + (1-\phi)\tau$, which is the marginal benefit of abatement, and equalises it with its marginal abatement costs $a y_0 \nu$. This rule is the channel through which reputation affects real outcomes. Holding the non-committed tax fixed, a higher ϕ raises the expected future carbon price whenever $\tau^c > \tau$, which increases abatement and lowers emissions.

3.2 Committed Government

The committed government c minimises the scaled net-present welfare

$$\int_0^\infty r e^{-rs}\, w(\tau_s^c; \phi_0, m_s)\,\mathrm{d}s\,, \tag{14}$$

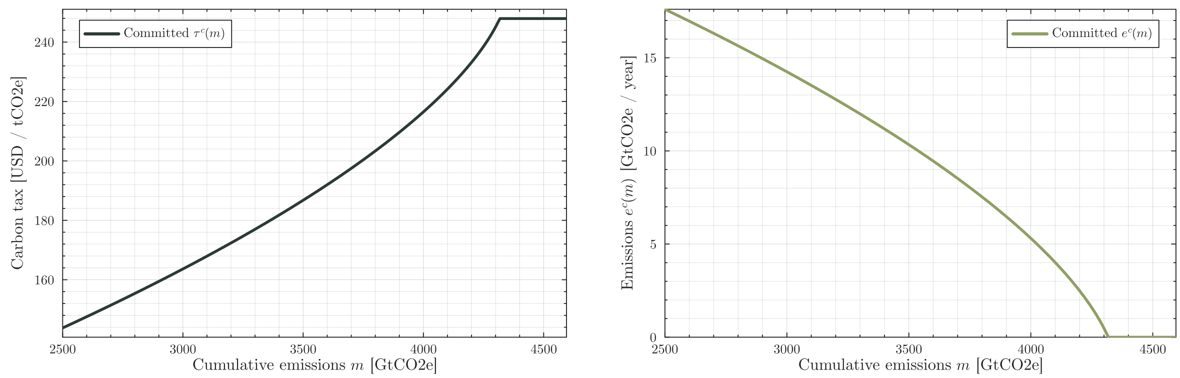

Figure 3: Optimal climate policy policy $\tau^c(m_t)$ and resulting emissions e under a certain committed government as a function of cumulative carbon concentration m_t

by picking a Stackelberg feedback tax schedule $\tau_t^c = \tau^c(m_t)$, subject to the cumulative emissions dynamics (5). Note that the beliefs of the firm ϕ_t when facing a committed government c do not change over time (11), and $\phi_t = \phi_0$ for all t.

Denote by $u^c(m_0)$ the minimised welfare costs (14). The committed policy is the reference path against which firms interpret future signals, and it also determines the high-reputation boundary of the non-committed problem. These welfare costs satisfy the HJB equation

$$ru^c(m) = \min_{\tau^c} \left\{ rw(\tau^c; \phi_0, m) + \big(e_0 - a(\phi_0, \tau^c)\big)\partial_m u^c(m) \right\}. \tag{15}$$

The optimal committed tax schedule τ^c satisfies the first condition

$$0 = r\left(\frac{\tau^c}{y_0\nu} + \frac{\partial l}{\partial \tau} + \frac{1}{y_0\nu}\frac{\partial l}{\partial a}\right) - \frac{1}{y_0\nu}\partial_m u^c(m). \tag{16}$$

When $\delta_R = \delta_A = 0$, the interior condition collapses to the Pigouvian benchmark $\tau^c = \partial_m u^c / r$. With transition losses, the committed tax also internalises the effect of taxation on stranded-asset costs, so it generally differs from a pure climate-damage tax.

Figure 3 shows the optimal committed policy $\tau^c(m_t)$ and the resulting emissions. A committed government sets a climate policy of approximately 150\$/tCO$_2$e at cumulative emission levels of $m_0 \approx 2500$GtCO$_2$e, which are consistent with cumulative emission levels in 2020. As cumulative emissions increases, the optimal committed carbon rapidly grows to 248\$/tCO$_2$e at approximately 4250GtCO$_2$e. At this tax level, firms abate all emissions, as 248\$/tCO$_2$e $\equiv \nu y_0 e_0$.

3.3 Non-Committed Government

The non-committed government nc picks, at each moment in time t, a tax τ_t to solve

$$u(\phi, m) = \inf_{\tau_s} \mathbb{E}_{\phi,m} \int_0^\infty re^{-rs} w(\tau_s; \phi_s, m_s) \, \mathrm{d}s \,, \tag{17}$$

subject to the belief dynamics (11) and cumulative emissions (5). Conditional on the current tax τ, the infinitesimal generator is

$$\mathcal{G}^\tau u = e(\phi, \tau) u_m(\phi, m) + \phi(1-\phi)\xi(\tau)^2 \left(-\phi u_\phi(\phi, m) + \frac{1}{2}\phi(1-\phi) u_{\phi\phi}(\phi, m) \right) . \tag{18}$$

Therefore, the Hamilton-Jacobi-Bellman equation is

$$ru(\phi, m) = \min_\tau \left\{ rw(m, \phi, \tau) + \mathcal{G}^\tau u(\phi, m) \right\} . \tag{19}$$

For more details on the solution of (19), see Appendix (B). The HJB highlights the reputation trade-off. A lower tax can reduce current transition costs, but it also changes firms' abatement through the expected tax and increases the distance from the committed benchmark. This distance enters $\xi(\tau)$ and determines how informative the signal is about the government's type. Hence, the value of reputation appears in the non-committed government's problem as an endogenous continuation cost of deviating from $\tau^c(m)$.

Figure (4) illustrates the optimal policy $\tau(\phi, m)$ as a function of its reputation ϕ and the cumulative emission m. For low levels of cumulative emission m, a non-committed government nc with high reputation ϕ leverages its reputation and sets a low climate policy, relying on the belief of firms in its committment. As cumulative emission m_t increase, this strategy becomes increasingly costly as climate damages $d(m_t)$ accumulate. Figure 5 shows the percentage deviation $\log(\tau(\phi, m)/\tau^c(m))$ between the committed and the non-committed tax. For low levels of cumulative emissions m_t, a non-committed government, regardless of the reputation, sets a lower climate policy than a committed government. As cumulative emissions increase, concerns of climate damages become more salient than concerns over transition losses (light red area). Then non-committed governments choose to "mimic" committed ones by setting the same climate policy $\tau(\phi, m) = \tau^c(m)$ (sand area). A non-committed government with high reputation, if cumulative emissions are high, can choose to temporarily strongly deviate from the committed climate policy $\tau(\phi, m) \ll \tau(m)$ (dark red area), to briefly capitalise on its reputation before choosing to mimic.

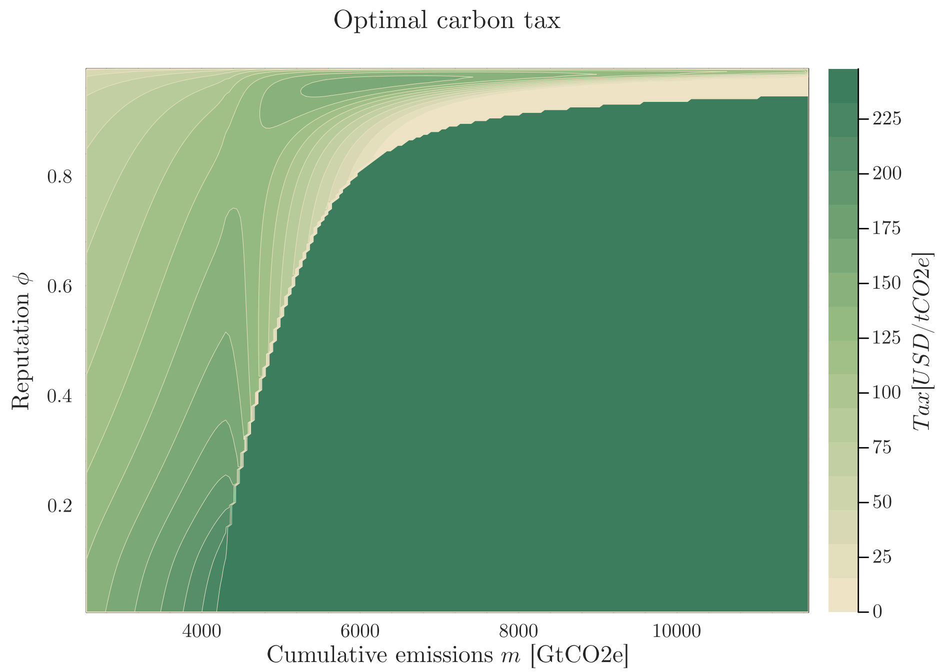

Figure 4: Equilibrium tax policy $\tau(\phi, m)$ for a non-committed government.

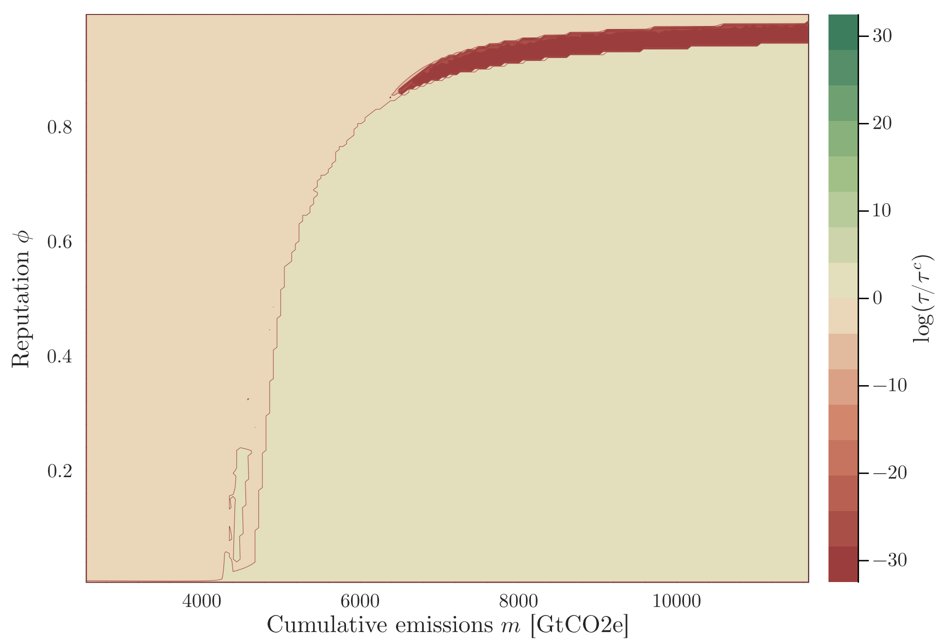

Figure 5: Percentage deviation $\log(\tau(\phi, m)/\tau^c(m))$ between the committed and the non-committed tax.

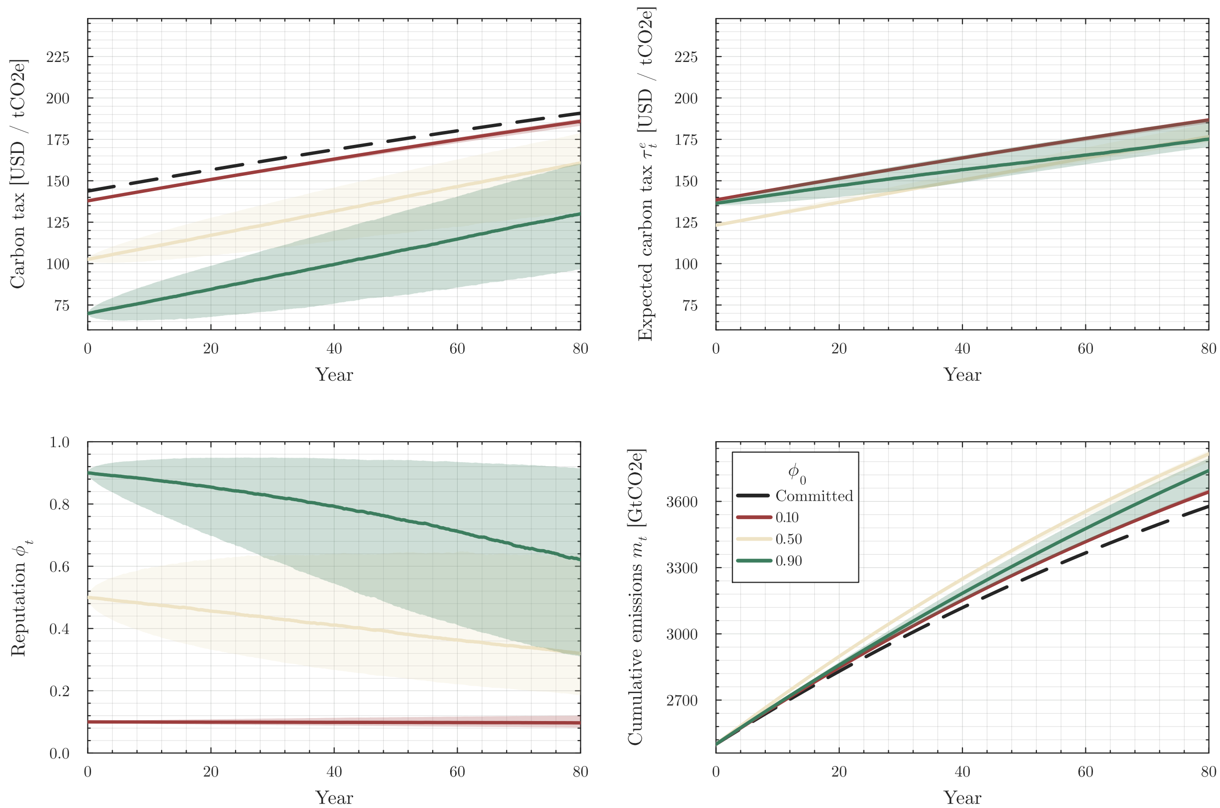

Figure 6: Dynamics of τ_t, $\tau^e = \tau^c\phi_t + (1-\phi_t)\tau_t$, ϕ_t, and m_t.

3.4 Dynamics

These policies imply dynamics for cumulative emissions m_t and reputation ϕ_t. Figure 6 illustrates these dynamics in a low, a middle, and a high initial reputation: $\phi_0 \in \{0.1, 0.5, 0.9\}$. The top left panel shows the path of the climate policy policy $\tau - t$, the top right panel shows the path of the expected climate policy policy from the point of view of firms $\tau^e = \tau^c\phi_t + (1-\phi_t)\tau_t$, the bottom left panel shows the path of reputation ϕ_t, and the bottom right panel shows the path of cumulative emissions m_t. The low reputation government does not have reputational weight to reduce the climate policy, without significantly lowering abatement investment by firms. Hence, they choose a climate policy path τ_t that closely mimics the committed tax path. Similarly, the higher the reputation, the lower the implemented tax, as the reputation can be expended to sustain firms abatement investments. Over time, these differences shrink as the reputation of all three types gets expended (bottom-left panel). Importantly, all three types manage expectations of firms and induce similar beliefs about climate policy (top-right panel).

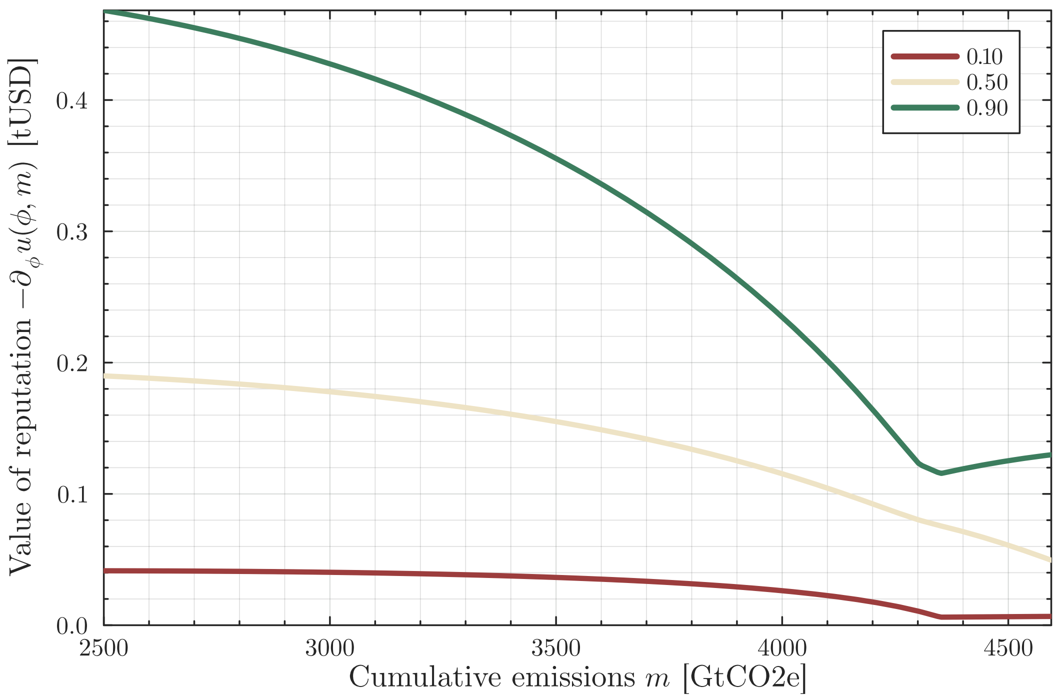

Figure 7: Value of reputation $-\partial_\phi u$ at different levels of reputation ϕ and cumulative emission levels m.

4 Value of Reputation

This section illustrates to important implication of the model. First, it quantifies the value of reputation and, as a consequence, potential losses from a reputational shock, such as the election of a government sceptical about the green transition. Second, it discusses the consequences of a positive shock to the noise σ. This can be interpreted as a period where an unanticipated war or energy crises affects the volatility in the ETS futures' market.

4.1 The Value of Reputation

We define the value of reputation as $-\partial_\phi u$, that is, the shadow value of reputation in lowering the net-present transition costs u. Figure 7 shows the value of reputation at different levels of reputation ϕ and cumulative emission levels m. At lower levels of cumulative emissions m reputation is more valuable, as the flow climate damages are not as high and, hence, the benefits of deviating from the committed climate policy path. If cumulative emissions are high, climate damages dominate the welfare costs of the transition and there is no government can credibly deviate from a high climate policy. Furthermore, the value of reputation is larger

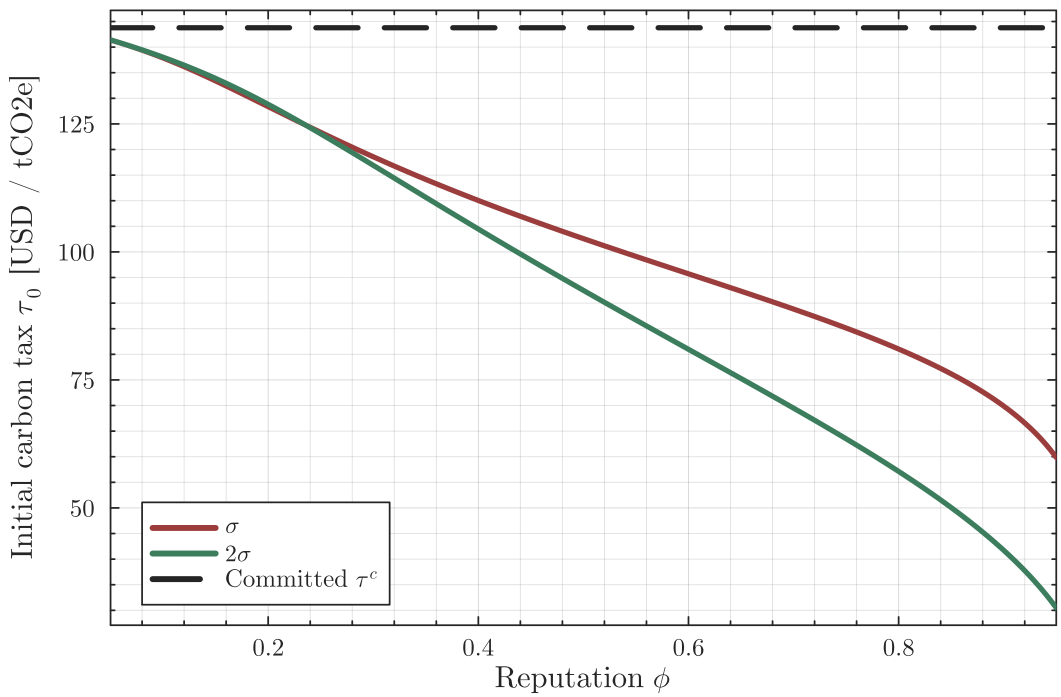

Figure 8: Initial climate policy τ_0 in a low σ and a high 2σ noise scenario.

for larger level of reputation, as a high reputation government can choose to systematically deviate, lower transition costs, and at the same time achieve high abatement levels.

This result suggests that governments with high reputation, especially early in the transition, can "get away" with large decreases in carbon taxation without lowering abatement from firms. This privilege, which gets eroded over time, is valued at around 500bn$. A government that loses this reputation, possibly due to the election of a climate sceptic party, loses this freedom and cannot adjust carbon taxes without compromising the green transition.

4.2 Periods of High Volatility

Now we consider the effect of a doubling of the signals' noise σ on the non-committed governments' policy τ. Figure 8 shows the initial climate policy τ_0 as a function of beliefs in a low σ and a high 2σ noise scenario. An increase in the noise makes it harder for firms to distinguish a committed from a non-committed government. In this situation, for a government with high reputation, it becomes easier to deviate further and lower climate policies. Going back to the interpretation of the increase in noise as an increase in the volatility of permits' future markets due to an energy crisis, this result rationalises the fact that it is easier for governments with high reputation to adjust policy in periods of crisis, without compromising their reputation and the green transition.

5 Conclusion

This paper studies climate policy when governments cannot fully commit to future carbon prices and firms learn about policy commitment from noisy signals. In this framework, reputation affects both firms' expected climate policy and the policy chosen by a non-committed government.

The results show that reputation provides the government with policy flexibility. A government with a strong reputation can set a lower current climate policy while firms continue to place substantial weight on the committed policy path. This allows the government to reduce current transition costs without generating an equally large decline in the expected carbon price and abatement. This flexibility is constrained by learning, since deviations from the committed path reveal information about the government's type and tend to erode its reputation over time. When policy signals are noisier, this reputational constraint is weaker and the government can deviate further.

Because abatement affects cumulative emissions, different reputation states can also generate different subsequent climate paths. The environmental effect of reputation therefore results from the interaction between firms' beliefs and the government's endogenous policy response. Overall, the analysis shows that beliefs about future policy constitute an additional state variable shaping current climate policy and its consequences.

References

Acemoglu, D., & Rafey, W. (2023). Mirage on the horizon: Geoengineering and carbon taxation without commitment. *Journal of Public Economics*, *219*, 104802. https://doi.org/10.1016/j.jpubeco.2022.104802

Baldwin, E., Cai, Y., & Kuralbayeva, K. (2020). To build or not to build? Capital stocks and climate policy. *Journal of Environmental Economics and Management*, *100*, 102235. https://doi.org/10.1016/j.jeem.2019.05.001

Barro, R. J., & Gordon, D. B. (1983). A Positive Theory of Monetary Policy in a Natural Rate Model. *Journal of Political Economy*, *91*(4), 589–610. Retrieved March 27, 2026, from https://www.jstor.org/stable/1831069

Brunner, S., Flachsland, C., & Marschinski, R. (2012). Credible commitment in carbon policy. *Climate Policy*, *12*(2), 255–271. https://doi.org/10.1080/14693062.2011.582327

Canzoneri, M. B. (1985). Monetary Policy Games and the Role of Private Information. *The American Economic Review*, *75*(5), 1056–1070. Retrieved March 27, 2026, from https://www.jstor.org/stable/1818645

Chiappinelli, O., & May, N. (2022). Too good to be true? Time-inconsistent renewable energy policies. *Energy Economics*, *112*, 106102. https://doi.org/10.1016/j.eneco.2022.106102

Cukierman, A., & Meltzer, A. H. (1986). A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information. *Econometrica*, *54*(5), 1099–1128. https://doi.org/10.2307/1912324

Dietz, S., & Venmans, F. (2019). Cumulative carbon emissions and economic policy: In search of general principles. *Journal of Environmental Economics and Management*, *96*, 108–129. https://doi.org/10.1016/j.jeem.2019.04.003

Faingold, E., & Sannikov, Y. (2011). Reputation in Continuous-Time Games. *Econometrica*, *79*(3), 773–876. https://doi.org/10.3982/ECTA7377

Faust, J., & Svensson, L. E. (2001). Transparency and credibility: Monetary policy with unobservable goals. *International Economic Review*, *42*(2), 369–397.

Gersbach, H., & Glazer, A. (1999). Markets and Regulatory Hold-Up Problems. *Journal of Environmental Economics and Management*, *37*(2), 151–164. https://doi.org/10.1006/jeem.1998.1063

Hansen, T. (2022). Stranded assets and reduced profits: Analyzing the economic underpinnings of the fossil fuel industry's resistance to climate stabilization. *Renewable and Sustainable Energy Reviews*, *158*, 112144. https://doi.org/10.1016/j.rser.2022.112144

Helm, D., Hepburn, C., & Mash, R. (2003). Credible Carbon Policy. *Oxford Review of Economic Policy*, *19*(3), 438–450. https://doi.org/10.1093/oxrep/19.3.438

Kalkuhl, M., Steckel, J. C., & Edenhofer, O. (2020). All or nothing: Climate policy when assets can become stranded. *Journal of Environmental Economics and Management*, *100*, 102214. https://doi.org/10.1016/j.jeem.2019.01.012

King, R. G., & Lu, Y. K. (2022). *Evolving reputation for commitment: The rise, fall and stabilization of us inflation*. National Bureau of Economic Research.

Kreps, D. M., & Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory*, *27*(2), 253–279. https://doi.org/10.1016/0022-0531(82)90030-8

Kydland, F. E., & Prescott, E. C. (1977). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, *85*(3), 473–491. Retrieved January 27, 2026, from https://www.jstor.org/stable/1830193

Lu, Y. K., King, R. G., & Pasten, E. (2016). Optimal reputation building in the New Keynesian model. *Journal of Monetary Economics*, *84*, 233–249.

Requate, T., & Unold, W. (2001). On the Incentives Created by Policy Instruments to Adopt Advanced Abatement Technology if Firms are Asymmetric. *Journal of Institutional and Theoretical Economics*, *157*(4), 536–554. https://doi.org/10.1628/0932456012974468

Semieniuk, G., Holden, P. B., Mercure, J.-F., Salas, P., Pollitt, H., Jobson, K., Vercoulen, P., Chewpreecha, U., Edwards, N. R., & Viñuales, J. E. (2022). Stranded fossil-fuel assets translate to major losses for investors in advanced economies. *Nature Climate Change*, *12*(6), 532–538. https://doi.org/10.1038/s41558-022-01356-y

Van Der Ploeg, F., & Rezai, A. (2020). Stranded Assets in the Transition to a Carbon-Free Economy. *Annual Review of Resource Economics*, *12*(1), 281–298. https://doi.org/10.1146/annurev-resource-110519-040938

A Calibration

A.1 Transition losses

The two transition-loss parameters are disciplined by mapping them into loss shares at economically distinct benchmark points. Let

$$\rho_R(a,\tau) := \frac{R(a,\tau)}{y_0} = \frac{\tau\sqrt{e_0(e_0-a)}}{y_0}, \qquad \rho_A(a,\tau) := \frac{A(a,\tau)}{y_0} = \frac{\tau(a-\omega e_0)_+}{y_0}. \tag{20}$$

δ_R and δ_A are large when a given exposure share generates a large real transition loss rather than only a fiscal transfer or a private redistribution.

In the numerical implementation, we store the calibration as two annual benchmark loss shares. The residual-exposure share is evaluated at observed abatement a_0 and carbon tax τ_0:

$$\ell_{R,0} := \frac{\delta_R}{2}\rho_R(a_0,\tau_0)^2, \qquad \delta_R = \frac{2\ell_{R,0}}{\rho_R(a_0,\tau_0)^2}. \tag{21}$$

With the calibration used here, the baseline $\delta_R = 10$ corresponds to $\ell_{R,0} \simeq 4.85 \times 10^{-4}$, or about 0.049 percent of output per year at (a_0, τ_0). The accelerated-retirement share is instead evaluated at the net-zero benchmark. Let $\tau^{nz} := y_0 \nu e_0$ be the tax that induces $a = e_0$ in the firm problem. Then

$$\ell_{A,0} := \frac{\delta_A}{2}\rho_A(e_0,\tau^{nz})^2, \qquad \delta_A = \frac{2\ell_{A,0}}{\rho_A(e_0,\tau^{nz})^2}. \tag{22}$$

Evaluating the inherited exposure before abatement gives

$$\rho_R(0,\tau) \simeq \frac{\tau e_0}{y_0} \in [0.017, 0.034]. \tag{23}$$

If the marginal excess transition burden is between 10 and 30 cents per dollar of carbon liability, then $\delta_R = m/\rho_R$ gives $\delta_R \in [3, 18]$ over this price range.

For the residual channel, I keep $\delta_R = 10$ as the baseline and $\delta_R \in [3, 30]$ as a conservative sensitivity range. For the accelerated-retirement channel, stranded-asset estimates are more naturally interpreted as stock losses along a transition path. Semieniuk et al. (2022) estimate about \$1.4 trillion in stranded upstream oil and gas assets under a medium expectations realignment, while Hansen (2022) estimates \$13–\$17 trillion in reserve devaluation under 1.5–1.8°C stabilisation scenarios. Annualising these stock losses over 15 years at $r = 0.03$ gives roughly \$0.12–\$1.42 trillion per year. Relative to the output normalisation in the code, this corresponds approximately to $\ell_{A,0} \in [6 \times 10^{-4}, 7 \times 10^{-3}]$. I therefore use $\ell_{A,0} = 5 \times 10^{-4}$

as a conservative baseline and $\ell_{A,0} \in [10^{-4}, 10^{-2}]$ as a sensitivity range.

A.2 Signal volatility

We calibrate the volatility of the policy signal using daily settlement prices of European Union Allowance (EUA) futures contracts, retrieved from the London Stock Exchange Group (LSEG) database. The sample period spans from the 29th of December 2010 to the 15th of June 2026, covering a total of 3931 trading days.

Let y_t^c denote the log price of contract c observed on trading day t. We assume that futures prices share a common latent component s_t, while allowing for contract-specific fixed effects and a linear time-to-maturity effect:

$$y_t^c = s_t + \alpha_c + \gamma d_{c,t} + u_t^c, \tag{24}$$

where α_c is a contract fixed effect, $d_{c,t}$ denotes the remaining time to maturity of contract c at date t, and u_t^c captures contract-specific measurement and microstructure noise. The time-to-maturity variable is expressed in years.

First-stage cross-sectional adjustment. Let C_t denote the set of contracts observed at date t, and let $N_t \equiv |C_t|$. We define the date-specific cross-sectional averages

$$\bar{y}_t = \frac{1}{N_t} \sum_{c \in C_t} y_t^c, \qquad \bar{d}_t = \frac{1}{N_t} \sum_{c \in C_t} d_{c,t}. \tag{25}$$

Taking cross-sectional deviations from these averages eliminates the common state s_t:

$$\begin{aligned} \widetilde{y}_t^c &\equiv y_t^c - \bar{y}_t \\ &= \alpha_c - \bar{\alpha}_t + \gamma \left(d_{c,t} - \bar{d}_t\right) + u_t^c - \bar{u}_t, \end{aligned} \tag{26}$$

where

$$\bar{\alpha}_t = \frac{1}{N_t} \sum_{c \in C_t} \alpha_c, \qquad \bar{u}_t = \frac{1}{N_t} \sum_{c \in C_t} u_t^c. \tag{27}$$

We use Eq. (26) to estimate the maturity slope γ together with a set of relative contract effects. Let $\widehat{a}_c$ denote the contract-specific component recovered from this first-stage regression. We then purge each observed futures price of the estimated maturity component and contract heterogeneity:

$$y_t^{c,\text{pure}} = y_t^c - \widehat{a}_c - \widehat{\gamma} d_{c,t}. \tag{28}$$

The resulting series provide noisy contract-specific measurements of a common latent component, up to a normalization in levels.

Latent-state interpretation. The first-stage transformation removes the cross-sectional maturity structure and contract-level heterogeneity before estimating the latent common component. Since the level of the latent state is not separately identified from the average contract effect, the Kalman filter recovers a normalized state, denoted s_t^*. This normalization is immaterial for the calibration of the signal volatility, which depends on innovations in the latent state rather than on its level.

When the composition of the observed cross-section is constant over time, the normalization is simply a constant shift:

$$s_t^* = s_t + \kappa, \qquad \Delta s_t^* = \Delta s_t. \tag{29}$$

Annualized state-space specification. The latent common component is estimated using a Kalman filter. Since the data are observed at the daily trading frequency, we set

$$\Delta t = \frac{1}{252}, \tag{30}$$

so that one unit of model time corresponds to one year with 252 trading days. The discrete-time state equations are

$$s_t^* = s_{t-1}^* + \mu_{t-1}\Delta t + \sigma\sqrt{\Delta t}\,\eta_t, \tag{31}$$

$$\mu_t = \mu_{t-1} + \sigma_\mu\sqrt{\Delta t}\,\zeta_t, \tag{32}$$

where η_t and ζ_t are independent standard Gaussian innovations. The corresponding measurement equation for the purged futures prices is

$$y_t^{c,\text{pure}} = s_t^* + \varepsilon_t^c, \qquad \varepsilon_t^c \sim \mathcal{N}(0, \sigma_u^2). \tag{33}$$

Under this formulation, the daily innovation variance of the latent signal is $\sigma^2\Delta t$, while σ is directly expressed as an annualized volatility.

Estimating the model by maximum likelihood over 3,609 trading dates and 113 distinct contracts yields

$$\widehat{\sigma} \simeq 0.38. \tag{34}$$

We therefore set the signal volatility in the quantitative model to $\sigma = 0.38$.[4]

[4]As a robustness check, we estimate the measurement equation in levels while retaining all 113 contract

B Numerical Solver

B.1 Lower Boundary

At the lower boundary, firms can commit to not abate $a = 0$ as they know the government to be strategic. Hence, the welfare costs satisfy

$$\begin{aligned}\underline{u}(m_0) &= r \int_0^\infty e^{-rs} y_0 d(e_0 s + m_0) \\ &= y_0 r \int_0^\infty e^{-rs} \left(1 - e^{-\frac{\gamma}{2}\zeta^2 (e_0 s + m_0)^2}\right) \mathrm{d}s \\ &= y_0 \left(1 - r \int_0^\infty e^{-\frac{\gamma}{2}\zeta^2 e_0^2 s^2 - (\gamma\zeta^2 e_0 m_0 + r)s - \frac{\gamma}{2}\zeta^2 m_0^2} \mathrm{d}s\right)\end{aligned} \tag{35}$$

B.2 Upper Boundary

At the upper boundary $\phi_t = 1$ for all t. Hence, the upper bound welfare is given by

$$\bar{u}(m_0) = \int_0^\infty r e^{-rs} w(0; 1, m_s) \, \mathrm{d}s \, . \tag{36}$$

This satisfies the equation

$$r\bar{u}(m) = rw(0; 1, m) + \left(e_0 - \frac{\tau^c(m_s)}{y_0 \nu}\right) \partial_m \bar{u}(m) \tag{37}$$

for all m.

B.3 Interior

Let $\tau^*(\phi, m)$ denote the minimiser in (19), and write

$$a^*(\phi, m) \coloneqq a(\phi, \tau^*(\phi, m)), \tag{38}$$

$$w^*(\phi, m) \coloneqq w(m, \phi, \tau^*(\phi, m)), \tag{39}$$

$$\xi^*(\phi, m) \coloneqq \xi(\tau^*(\phi, m)). \tag{40}$$

fixed effects:

$$y_t^c = s_t + \alpha_c + \gamma d_{c,t} + u_t^c.$$

The full fixed-effect state-space specification delivers an annualized estimate of approximately $\widehat{\sigma} = 0.39$, compared with approximately $\widehat{\sigma} = 0.38$ under the two-step cross-sectional procedure. The similarity of the two estimates indicates that the calibration is not materially driven by the treatment of contract fixed effects or by changes in the composition of C_t.

In the interior, for $\phi \in (0,1)$, the value function satisfies

$$
\begin{aligned}
ru(\phi,m) = rw^*(\phi,m) + (e_0 - a^*(\phi,m))\,\partial_m u(\phi,m) + \phi(1-\phi)(\xi^*(\phi,m))^2 \\
\times\left(-\phi\partial_\phi u(\phi,m) + \frac{1}{2}\phi(1-\phi)\partial_{\phi\phi}u(\phi,m)\right).
\end{aligned}
\tag{41}
$$

The policy $\tau^*(\phi,m)$ is computed pointwise from

$$
\begin{aligned}
\tau^*(\phi,m) \in \arg\min_{\tau\geq 0}\Bigg\{ & rw(m,\phi,\tau) + (e_0 - a(\phi,\tau))\,\partial_m u(\phi,m) \\
& + \phi(1-\phi)\xi(\tau)^2\left(-\phi\partial_\phi u(\phi,m) + \frac{1}{2}\phi(1-\phi)\partial_{\phi\phi}u(\phi,m)\right)\Bigg\}.
\end{aligned}
\tag{42}
$$

When $a(\phi,\tau) < e_0$, an interior optimum satisfies

$$
\begin{aligned}
0 = r\left[\frac{1-\phi}{y_0\nu}\left(y_0 c'(a) + \partial_a l(a,\tau)\right) + \partial_\tau l(a,\tau)\right] - \frac{1-\phi}{y_0\nu}\partial_m u(\phi,m) \\
- 2\phi(1-\phi)\frac{\epsilon}{\sigma}\xi(\tau)\left(-\phi\partial_\phi u(\phi,m) + \frac{1}{2}\phi(1-\phi)\partial_{\phi\phi}u(\phi,m)\right),
\end{aligned}
\tag{43}
$$

evaluated at $\tau = \tau^*(\phi,m)$ and $a = a^*(\phi,m)$.