The Price of Delay: infrastructure lock-in and the carbon tax lock-in premium in the oil sector

Léo Jean*

June 15th, 2026

PRELIMINARY DRAFT - PLEASE DO NOT CITE OR CIRCULATE

Abstract

Delaying climate policy does more than postpone abatement: it encourages fossil fuel investments whose emissions lock in once capital is sunk. I quantify this using field-level data on around 30,000 oil assets. A 1.5°C-aligned carbon tax delayed from 2016 to 2030 overshoots the oil carbon budget by 33% (19% for an increasing tax). A third stems from oil fields developed during the delay rather than from overproduction. I introduce the carbon tax lock-in premium: the additional tax needed after a delay to offset committed emissions from investments that would not have occurred under timely policy. Though modest relative to the tax itself, ignoring it leaves a residual overshoot of 17–19 GtCOeq, over 10% of the remaining 2030 budget. This cost builds fast: even five years of locked-in investment add over 10 GtCOeq. Because lock-in concentrates in a few deepwater projects, targeted supply-side measures can substantially cut the cost of delay.

Keywords: Climate change, fossil fuels, oil, carbon mitigation, carbon tax, climate policy delay, carbon lock-in, infrastructure lock-in, climate policy credibility.

JEL Codes: Q35, Q52, Q54, Q58, L71.

*Paris School of Economics, ENS - PSL. Email: leo.jean@psemail.eu.

1 Introduction

Renewed states commitments towards lowering carbon emissions lead carbon-intensive companies to expect future ambitious climate policies, especially in the oil and gas sector. In 2015, the Paris Agreements set a common objective for states parties to hold the rise in global temperatures to "well below 2°C [...] and pursuing efforts to limit [it] to 1.5°C". In order to remain within these limits, significant portions of existing fossil fuel reserves must remain unextracted (McGlade and Ekins, 2015; Welsby et al., 2021). To achieve this goal, the economic literature widely recognises carbon taxation as the most effective tool for curbing fossil fuel use in the coming century. However, despite this recognition, as of 2023, only a few regions or countries worldwide have implemented such taxes, and their levels fall far short of what is required to align with the objectives outlined in the Paris Agreement framework (Metcalf, 2019).

This lack of stringency in current climate policies leads to two consequences. First, it results in overproduction from existing production facilities, as studies have shown that low carbon taxes exert minimal influence on current oil and gas production (Heal and Schlenker, 2019). Second, it does not act as a sufficient deterrent to change producers' investment strategies. A large share of oil and gas firms use carbon shadow pricing in order to assess their investments to account for future climate policies (Fried et al., 2021). However, Ahluwalia (2017) shows that these shadow prices are usually low, although higher than observed carbon prices, ranging from \$23 to \$80 per metric ton[1], leading oil companies to only marginally adapt their investments to climate change. Low observed shadow prices used by oil companies highlight the poor credibility of States when they commit, and have long-lasting impacts on the path of oil investments and supply. The long term impact of new investments in fossil fuels is due to the inertia of carbon emissions associated with it and is referred to as carbon lock-in. Once investments are made into new producing infrastructure, carbon emissions from the associated deposits can become locked-in for long time periods, and render a low-carbon energy transition more difficult. Failing to account for these dynamics can lead to underestimate the real cost of delaying climate action, and more importantly underestimate the level of carbon tax required in the future in order to meet climate objectives.

In this paper, I use detailed field-level data on oil production facilities to study the

[1]Prices higher than \$60 are only used to stress test some highly-exposed investments

long-term consequences of delaying carbon taxation in the oil production sector from 2016 to 2030. I show that policy delay leads to a substantial overshoot of the carbon budget set in 2016, driven not only by excess production during the delay but also by long-lived investment decisions that lock in future emissions. Crucially, I show that these investment dynamics create a wedge between a naïve tax, computed under the assumption of fixed extraction costs, and the optimal tax required after the delay. I call this gap the carbon tax lock-in premium. Ignoring lock-in thus leads to systematically underestimating carbon prices and translates into additional post-tax emissions overshoot. Finally, I identify the new oil projects most responsible for these long-term costs and assess how targeted supply-side policies can mitigate both overproduction and lock-in.

Climate policy delays and the feasibility of climate objectives have been extensively studied within the integrated assessment models (IAMs) literature. Most studies conclude that delaying ambitious climate policies increases substantially mitigation costs, but a delay up to 2030 still allows for the 450ppm or 2°C objectives to be attained (Eom et al., 2015; Riahi et al., 2015; Kotlikoff et al., 2021). Making up for the inaction until 2030 renders strong and rapid adaptation necessary in the 2030-2050 period. However, as noted by Van Vuuren et al. (2016), the feasibility of the expected decarbonisation rates in this period is far from assured, and the potential of carbon capture and storage (CCS) might be overestimated in these models. A strong focus is placed on the green paradox[2], while carbon lock-in effects and path dependency remain understudied. Studying a large variety of IAMs, Grubb et al. (2021) show that the majority only partially takes into account the dynamic nature of investments and production processes, and considers abatement costs in each period to be independent from prior abatement. Path dependency, inertia, and induced innovation are largely disregarded.

Among these models, only IMACLIM-R, REMIND, and WITCH incorporate these three features, and can thus result in unattainable climate objectives following delays in climate action. In Riahi et al. (2015)' study, only the IMACLIM-R model leads to the conclusion that the transition is infeasible with a delay up to 2030. Similarly, using these three models, Jakob et al. (2012) find that delaying the implementation of ambitious climate

[2]Coined by Sinn (2008) and extensively studied afterwards (see e.g. van der Ploeg and Withagen (2015)), the green paradox states that a delay in implementing carbon policy can encourage fossil fuels producers to overproduce during the delay (as compared to a "no policy" scenario) instead of losing their resource after the policy is implemented.

policies until 2030 makes it impossible to attain the 450ppm goal. Luderer et al. (2013) and Bertram et al. (2015) find similar results using the same models. Assumptions on the lifetime of infrastructures are crucial. Carbon lock-in in these studies is probably overstated because the models assume that infrastructure built is used to the end of its lifetime. Concurrently, Insley (2017) shows that early retirement of carbon-intensive infrastructure is to be expected if carbon policy is implemented. Nevertheless, despite capturing some carbon lock-in, the complexity of these models limits the understanding of the mechanisms underlying lock-in effects and the identification of specific investments leading to it (Grubb et al., 2021). Additionally, IAMs' modelling of fossil fuel supply often produces unrealistic responses to changes in demand (Höök and Tang, 2013). Despite accounting for heterogeneity across fields or regions, production costs are often considered constant at the field level, and computed as the breakeven price of the project (as in Heal and Schlenker (2019)). Investors' logic and geophysical constraints are largely disregarded, leading to strong adaptation in the allocation of the supply. In reality, when choosing to develop a field, investors determine the optimal level of infrastructure needed, which in turn determines the ultimate recovery rate of the field and its production rates (Venables, 2014). Production rates once the field has started its production cannot be adjusted at will and are mainly dictated by geophysical and technological constraints, to the point that no adaptation to changes in oil price can be observed at the well level (Anderson et al., 2018).

An increasing body of literature has shifted its focus towards carbon lock-in. Carbon lock-in refers to the phenomenon whereby investments in carbon-intensive infrastructures and technologies induce long-term dependence on high carbon emissions and hinder the transition to low-carbon alternatives. Developed by Unruh (2000), this concept includes three main dimensions identified by Seto et al. (2016) that are mutually intertwined and contribute to the inertia of carbon emissions: infrastructural, institutional and behavioural. This study focuses on the first dimension, through which newly developed oil deposits and associated infrastructures can render more difficult the phasing-out of oil supply and associated emissions. Studying carbon lock-in in this setting implies a change in perspective from annual emissions to committed emissions over the future associated with each infrastructure (see Davis and Socolow (2014) for more details on the commitment accounting of carbon emissions). Several papers tried to estimate the quantity of carbon emissions embedded in existing carbon-intensive infrastructures. Unsurprisingly, the later a study is published,

the more embedded emissions are found to jeopardise climate change mitigation objectives (Smith et al., 2019; Tong et al., 2019), with Trout et al. (2022) finding that current fossil fuels producing infrastructure would warm the world beyond 1.5°C.

Additional papers adopt a dynamic perspective to assess carbon lock-in from carbon-intensive infrastructures (see Fisch-Romito et al. (2021) for a review of this literature). Lazarus et al. (2015) offers an initial assessment, though partial, of the carbon lock-in arising from fossil fuel supply infrastructure. Wachtmeister and Höök (2020) compare the investment and production cycles of conventional and unconventional tight oil to assess the consequences of delayed climate action[3]. In both these studies, the underlying logic is that once capital investments are incurred, the decision to produce or not does not depend on initial expenses. The decision to invest relies on the expected net present value of the project, whereas the decision to produce only depends on extraction costs. Oil projects are on average costly to develop but cheap to operate (Lazarus et al., 2015; Venables, 2014), making newly developed projects a burden for future climate action. Fossil fuel supply results from dynamic, discrete and long-lasting investment choices which determine a path of production for every project. In this sense, newly developed projects represent committed emissions over the future, and lead to a carbon-intensive path dependency. Consequently, the academic literature has seen a rise in support for supply-side policies (Erickson et al., 2018; Green and Denniss, 2018; Fæhn et al., 2017), with an emphasis put on banning new oil projects (Green et al., 2024; Erickson and Lazarus, 2018; Van Asselt, 2021). Addressing the supply is more and more perceived as complementary to demand-side policies. On the one hand, it provides numerous advantages in terms of feasibility: low administrative costs, higher abatement certainty, higher public support. On the other hand, it allows to mitigate the cost arising from carbon lock-in risks (Green et al., 2024). Nevertheless, as noted by Lazarus and van Asselt (2018), the literature on demand-side policies far outweighs the one on supply-side policies, despite its apparent benefits.

To date, no study has offered precise estimates of the emissions consequences resulting from carbon lock-in in the oil production sector, taking into account the dynamic nature of investments. Additionally, while the literature on carbon lock-in emphasizes the persistence of emissions due to long-lived infrastructure, it has not translated these effects into impli-

[3]The characteristics of each source of oil has tremendous impacts on the associated dynamics of investment, and affect in turn the potential lock-in effects emerging from it.

cations for the level of carbon taxation following policy delays. This paper introduces the concept of a carbon tax lock-in premium, defined as the additional carbon price required to offset committed emissions resulting from delayed policy implementation and associated over-investments.

This paper aims at filling these gaps in the literature by investigating the consequences of a 14 years delay, between 2016 and 2030, in the implementation of carbon taxation in the oil producing sector. Both a constant carbon tax and a carbon tax increasing at a 2% discount rate are assessed. I find that for a 100$/tCO2eq constant carbon tax (respectively 57$/tCO2eq increasing one), in line with the updated oil carbon budget set by Welsby et al. (2021), the delay induces a 33% overshoot of the budget (respectively 19%). Two thirds of the overshoot occurs at the intensive margin, the rest originating from fields developed during the delay. From a carbon lock-in perspective, results suggest that overcommitted emissions represent more than half of total committed emissions during the period, and are concentrated in a small number of specific oil projects (mainly deepwater offshore). Overcommitted emissions from these assets appear resilient to higher taxes. As a result, achieving the same carbon budget after a policy delay requires a systematically higher carbon tax than suggested by a naïve calibration ignoring investment-driven lock-in. Failing to take into account this lock-in effect leads to a residual overshoot of the budget ranging from 15 GtCO2eq to 20 GtCO2eq, more than 10% of the remaining budget. I refer to this additional required tax as the carbon tax lock-in premium, and assess its implications for policy stringency over time as lock-in builds up.

The remainder of the paper is organized as follows. Section 2 presents the data and methodology used to estimate the magnitude of the carbon budget overshoot and to identify investments responsible for carbon lock-in effects. Section 3 presents results for (i) yearly overemissions stemming from the delay, (ii) the role played by carbon lock-in and (iii) the carbon tax lock-in premium required to compensate lock-in effects. Section 4 tests the sensitivity of my results to a large array of alternative assumptions, including the inclusion of time-varying tax pass-through from producers to consumers. In section 5, the implications of the results are discussed and policy recommendations are provided to mitigate the cost of carbon lock-in. Finally, Section 6 concludes.

2 Data and methods

2.1 Oil and gas field data

This study relies primarily on proprietary micro-level data from Rystad Energy. Rystad's UCube database[4] covers the vast majority of oil and gas production in existing fields worldwide since 1900. It contains as well informations on fields that are expected to be developed and discovered up to 2100. The analysis is conducted at the "asset" level, the smallest geographic scale available in the database, which typically corresponds to a field or to a portion of a field. The data comprises more than 100,000 oil and gas producing assets, among which 20,322 are producing in 2025, when data was extracted. For each of these assets, the data set contains a breakdown of annual production by oil types and of costs by category, as well as various administrative, geographic and geological characteristics. Annual costs include operating expenses (disaggregated into production, transportation and SG&A OPEX), capital expenses (facility, well), exploration expenses, and government take (income taxes, royalties, taxes reported as OPEX, bonuses, etc)[5].

2.2 Field-level carbon intensities.

In addition to the Rystad database, I estimate lifecycle carbon intensities (CI) for each asset by combining upstream, midstream, and downstream emissions. This section provides an overview of the methods used to estimate these CI. Further details about the methods are described in depth in Appendix B.

For upstream CI, I use data obtained from Masnadi et al. (2018) for 737 oil fields, matched to 1,136 Rystad assets, for which I compute CI using the Oil Production Greenhouse Gas Emissions Estimator (OPGEE)[6]. A key driver of upstream CI variability is flaring intensity, for which information is scarcely available. I thus use estimates of flaring intensity derived from satellite observations (2012–2024) observed by NOAA's VIIRS instruments Elvidge et al. (2015), converted into GHG emissions using their algorithm. These observa-

[4]Information on this database is available at https://www.rystadenergy.com/services/upstream-solution.

[5]See appendix A for a description of each category of cost.

[6]The OPGEE model is openly accessible at: https://eao.stanford.edu/research-project/opgee-oil-production-greenhouse-gas-emissions-estimator

tions are geographically matched to Rystad assets following the procedure detailed in Ghosh et al. (2025). For assets without direct flaring data, I predict flaring-to-oil ratios (FOR) using a multi-stage matching procedure and assign reserve-weighted averages for unmatched observations. I then fit a regression model incorporating FOR, steam-to-oil ratio (SOR), gas-to-oil ratio (GOR), offshore status, and oil type to predict upstream CI for all remaining assets, including not-yet-developed ones.

For midstream and downstream CI, I use the Petroleum Refinery Life-Cycle Inventory Model (PRELIM)[7] and the Oil Products Emissions Module (OPEM)[8]. PRELIM estimates refining emissions based on crude assay characteristics, primarily API gravity and sulfur content (Abella and Bergerson, 2012), while OPEM accounts for transport and end-use emissions. I match 493 PRELIM crude assays publicly available within the two models to 392 Rystad crude streams (covering 64.4% of 2024 production) and predict CI for unmatched assets using API gravity, sulfur content, and refinery configuration (deep conversion, hydroskimming, or medium conversion). For assets lacking API gravity and/or sulfur content data (12% of 2016 reserves), I assign reserve-weighted averages from similar fields (e.g., same reservoir, sub-basin, or oil type).

Finally, I sum upstream, midstream, and downstream CI to generate a unique lifecycle CI estimate for each asset. This approach, combining direct OPGEE/PRELIM/OPEM computations, satellite-derived flaring data, and predictive modeling, ensures robust CI estimates across all assets, even those with incomplete data. In 2024, the production-weighted average carbon intensity is 466 kgCO2eq/bbl, with a weighted standard deviation of 112 kgCO2eq/bbl. Figure 17 reports cumulative reserves in 2016 ordered by lifecycle carbon intensity. A large share of reserves are situated around, with a small tail to the left with cleaner assets, and a larger tail to the right with dirtier assets. The graph splits cumulative reserves according to assets development timing. No significant difference exist between assets developed during the delay and assets already producing in 2016.

[7]The PRELIM model is openly accessible at: https://ucalgary.ca/energy-technology-assessment/open-source-models/prelim

[8]The OPEM model is openly accessible at: https://ociplus.rmi.org/methodology

2.3 Data selection and definition of costs

Assets selection. A selection on the data was performed so as to keep only oil-producing assets (gas is disregarded in this analysis). Among oil-producing assets, only those which reserves are discovered by 2025[9]. This leaves me with 28,327 assets with available oil reserves in 2016, among which 16,309 already producing in 2016.

Reserves and oil definition. Reserves used in this study corresponds to oil resources as defined by Rystad, i.e. all economically viable recoverable oil over the lifetime of the asset. Quantities are typically larger than 2P reserves because it includes reserve growth over time, generally due to geological uncertainty declining as production takes place. Oil is defined as any type of crude oil (bitumen, extra heavy, heavy, regular, light or sour), as well as condensate and natural gas liquids (NGL). The last two hydrocarbons are assimilated to oil in reason of their liquid nature and the fact that they are sold as light types of oil (and often mixed with other crude oils before being sold in order to reduce their viscosity). As explained in the next subsection, although NGL is sold at a much lower price than oil, its price is primarily dictated by oil price and not by gas price. I thus include it in my definition of oil. A given asset can produce only one type of crude oil, but can also produce condensate and/or NGL as well.

Development costs. Following Rystad's disaggregation of costs, only expenses imputed to oil operations are kept. Capital expenses related to wells drilling and the building of facilities during the pre-production period and in the early years of production correspond to the development cost of the assets. Following the definition given by Rystad, the threshold before which all capital expenses are assimilated to development ones is set at two years after production startup. For modelling purposes, development expenses occuring once production has started are gathered in the year preceding production startup (and discounted accordingly in order to leave projects NPV unchanged).

Production costs. As stated above, only production costs imputed to oil are taken into account. I define as production costs all operating expenses (production opex, trans-

[9]Since Rystad fully models production, costs, and reserves characteristics of not-yet discovered assets, I prefer to exclude them from the analysis. Sensitivity analysis includes a specification in which not-yet discovered assets are kept, and the associated exploration costs treated as development costs.

portation opex, and selling, general and administrative opex), as well as all non-development capital expenses. This implies that expenses directed at building new facilities and new wells after the development phase of the asset occured (i.e. later than two years after production startup) will be included in productions costs.

Government take. Government takes are disregarded in the analysis, to the exception of taxes reported as OPEX by companies (variable "Taxes Opex" in appendix A). This last category of taxes is reported as operating costs by producers, and are not directly linked to the companies revenues or any asset's profitability. I thus include it in production costs. Among other types of government takes, some are based on profits from an asset (income tax), or expensed before exploration occurs (bonuses). I assume that these taxes do not affect investors in their decision to develop or not an asset, since the former affects the magnitude of the asset net profitability (after carbon taxation if there is any) but not its sign, and the latter is already expensed before development decisions are made. Fiscal regimes based on royalties could render a project unprofitable since its tax base is gross revenues. However, I assume that in presence of a worldwide carbon tax, governments would renegotiate and adapt royalty rates in order to keep assets active and continue extracting a share of the rent. Nevertheless, Ahlvik et al. (2022) and Stefanski et al. (2025) show that oil taxes affect substantially investments and production decisions. Given the complex and interdependent role of fiscal regimes on investment decisions (Erickson et al., 2017; Ahlvik and Harding, 2025), this study focuses on market-driven costs and carbon pricing mechanisms, excluding government takes from the analysis.

2.4 Modelling framework

The model of oil extraction in this paper relies on a sequence of decisions made by investors at each period (year) for each asset. Decisions made for each asset are independent of other assets. At each period, investors thus decide to invest or not in the development of the asset for undeveloped ones, and to produce or not for an active one.

Investment decision. I assume investment decisions made by investors to be primarily dictated by the expected net present value (NPV) of the project. Even though discounted cash-flow analysis is not the only criterion used in the oil industry to determine which field to

develop (risks are also assessed regarding political, financial, and safety hazards), it remains the cornerstone of project-based decision-making (Bailey et al., 2000; Jahn et al., 2008). Moreover, as noted by Erickson et al. (2017) and Wood Mackenzie (2013), it is the method used by the oil industry itself in their analysis. I use a 10% discount rate to compute these NPVs, which is the most common discount rate used in the oil and gas sector (Bailey et al., 2000), and is the one used by Rystad to compute breakeven prices, i.e. the price of oil above which the NPV of a project is positive. Robustness checks include the use of a 5% and a 15% discount rates.

In T_0, the decision to invest or not in an asset i depends on the anticipations of investors regarding the introduction of a carbon tax τ, and is thus governed by the equation:

$$\sum_{t=T_1}^{T} \frac{P(1-d_i)q_{it} \ - c_i}{(1+\delta)^{t-T_0}} - C_{iT_0} - \sum_{t=T_\tau}^{T} \tau \ \theta_i \ q_{it} \frac{(1+r)^{t-2016}}{(1+\delta)^{t-T_0}} \geqslant \ 0 \tag{1}$$

where T, and T_τ stand respectively for the expected last year of production and anticipated year of tax implementation. The production of the asset q_{it} is sold at a constant individual price $P(1-d_i)$, which corresponds to a reference oil price minus a discount specific to the asset. The costs are divided between development costs C_{iT_0}, assumed to be expensed 1 year before production startup, and production costs c_{it}. The carbon tax paid is equal to the carbon intensity of the asset θ_i multiplied by an initial 2016 carbon tax τ increasing at rate r. Finally, future production revenues and costs are discounted at a rate δ equal to 10%.

Rearranging equation (1), the maximum anticipated level of tax for the asset to be developed is as follows:

$$\bar{\tau} \ = \ max_{\{T\}} \ \frac{P(1-d_i) \ - \ BEP_i^{Project}(T)}{\theta_i \ Q_i^{\tau}(T)} \tag{2}$$

with $BEP_i^{Project}(T) = \frac{C_{iT_0} + \sum_{t=T_1}^{T} c_{it} \ (1+\delta)^{-t}}{\sum_{t=T_1}^{T} q_{it} \ (1+\delta)^{-t}}$ the breakeven price of the project including development costs up to the year T, and $\theta_i \ Q_i^{\tau}(T) = \theta_i \ \frac{\sum_{t=T_\tau}^{T} \frac{q_{it} \ (1+r)^{t-2016}}{(1+\delta)^{t-T_0}}}{\sum_{t=T_1}^{T} \frac{q_{it}}{(1+\delta)^{t-T_0}}}$ the anticipated share of discounted emissions taxed in 2016 tax terms.

Investors choose T, the last year of production, in order to maximise the net present value of the asset. In the absence of tax, T is always the last year of production predicted by

Rystad. However, the set of anticipations affects this maximisation and can lead investors to expect an asset lifetime different from what will effectively be the case once production has started and uncertainty about the tax has resolved. $\bar{\tau}$ thus corresponds to the maximum carbon tax consistent with a non-negative NPV, given anticipated taxation and optimal production duration.

In some observed cases, development does not occur in one year, and can sometimes span over multiple years, or stop at some point and resume later. This last situation is common for assets which development started in 2019: after the breakout of the COVID-19 pandemic and the subsequent drop in oil prices, many developments were stopped, and before resuming years later. This has an impact on the development status of an asset in each scenario. Therefore, the maximum level of tax for the development of the asset to be economically viable is computed for *each* development stage, with C_{iT_0} defined as the *remaining* amount of development costs. The asset is developed if and only if the anticipated level of tax is below the maximum tax $\bar{\tau}$ at each development stage. This is straightforward for assets which are developed before the end of the delay, since $\bar{\tau}$ increases as C_{iT_0} gets lower. However, development expenses can occur and suddenly stop when the tax is implemented, as uncertainty resolves and $Q_i^{\tau}(T)$ increases to 1 (or more in the case of a increasing tax). The development costs already incurred for these assets are lost, and the corresponding infrastructures built are abandoned.

Production decision. Once an asset is developed, the quantities produced are determined by comparing production costs with the oil price minus the tax (which is equal to zero during the delay). Production costs are computed as breakeven prices by stage of production, and are strictly increasing in time. Rystad predicts long lifetime for assets because it depends on a minimisation of the breakeven price of the full project, including development expenses. When removing these expenses, the breakeven price of production can be minimum with a final year of production earlier than forecasted by Rystad. This last characteristic stems from the fact that (i) operating costs increase over time, and (ii) as the production rate declines, capital expenses per barrel increase. Therefore, I identify for each asset stages of production with strictly increasing breakeven prices. In order to do so, I compute the last year of production that minimises the breakeven price, which corresponds to the end of stage 1, and repeat the operation for remaining years to identify subsequent

stages. For assets developed after 2015, the average number of production stages is 5.8, and the average duration of a production stage is 5.2 years, stage 1 usually being the longest.

Investment and production paths. The stream of investments and production assessed for each asset is the one observed up to 2025 and predicted by Rystad for the future. The rationale behind this assumption is twofold. Regarding the shape of the path, production rates (plateau and decline) mainly depends on initial investments and cannot be changed at will afterwards because of geological and technological constraints (Anderson et al., 2018; Venables, 2014; Cairns and Smith, 2019). Changing the production path would thus entail modifying development expenses, as well as all subsequent costs. Regarding the timing of the decisions, (i) during the delay, carbon taxation is not or only partially anticipated by investors, making the "business-as-usual" scenario forecasted by Rystad relevant, and (ii) once the tax is implemented the year of development does not matter, since I focus on cumulative production up to 2100.

One could argue that investors anticipating a carbon tax to be implemented, *a fortiori* an increasing one, would be incentivised to push their production and investments earlier in time, giving rise to a Green paradox effect. This valid concern would imply higher investments in earlier years, i.e. during the delay, and higher cumulative production for a given level of tax. This means that this study provides a lower bound of the cost of delaying climate policy.

Price assumptions. The model only focuses on supply-side responses, and therefore ignores the effect on oil prices of the implementation of a carbon tax. The level of tax assessed in the paper can thus be interpreted as the level of tax borne by producers. Section 4 includes a sensitivity analysis of my results incorporating carbon tax initially partially borne by consumers and progressively passed on to producers. I use exogenous predicted evolutions of the pass-through rates by carbon tax increase rates drawn from Heal and Schlenker (2019).

The reference oil price (Brent) is assumed to be constant and equal to \$80/bbl. Two rationales justify this choice. First, changing the level of price would only change the optimal level of carbon tax in my computations, without significantly changing the results in terms of carbon budget overshoot. Only the overall pool of economically viable assets (assets with breakeven prices lower than the oil price) is affected by this assumption. Yet, only few

projects require prices higher than \$80/bbl to be developed. Robustness checks include an oil price of \$100/bbl and \$60/bbl. A second rationale is the coherence of a \$80/bbl price as compared to oil markets dynamics. The stability of the oil market is mainly guaranteed by the 3 core countries of OPEC, namely Saudi Arabia, Kuwait and United Arab Emirates (Kellogg, 2024). For these countries, the IMF reports average past and projected (in 2025) fiscal breakeven prices of \$80/bbl (past) and \$84/bbl (future) for Saudi Arabia, and below for the two other countries (International Monetary Fund, 2024). Fiscal breakeven prices corresponds to the required oil price for oil exporting countries to balance their budget. Given the influence exerted by core OPEC countries on oil price, using their expected fiscal breakeven price as a reference for future oil price seems appropriate. Overall, using a \$80/bbl oil price allows me to model and capture invested infrastructure decisions and lock-in effects without biasing results toward extreme low-price environments where investment incentives collapse.

One might argue that the introduction of a carbon tax would alter long-term oil price forecasts. While this is true in principle, I argue that Middle Eastern countries would be significantly less affected by a carbon tax than other regions, due to their relative cheap costs and low carbon intensity, even more so when it comes to OPEC-core countries. Should the carbon tax be levied by producer States, OPEC-core countries would thus see their revenues stemming from oil production only moderately reduced by the tax. Results in figure 19 support this rationale.

Price volatility is also disregarded in the analysis because investment decisions are mainly based on expected oil prices in the long-term, especially for long-lived investments which are the key focus of this study. This assumption is further motivated by results from the literature. Pierru et al. (2018) shows that OPEC's use of its spare capacities significantly reduces oil price volatility and reinforce trust from investors in a stable market. Saudi Arabia acts as a swing producer, keeping significant spare capacities in order to adjust its production level quickly in response to demand shocks (Fattouh, 2007). Additionally, Berntsen et al. (2018) show that effects of price volatility on oil and gas investments in Norway are insignificant, while Tang et al. (2017) argue from a theoretical setting that it does not affect the decision to invest or not but could rather induce earlier investments.

Asset-specific discounts. While oil prices generally refer to reference prices such as

Brent, WTI or WCS, each asset sells its oil at a unique price corresponding to a reference price minus a discount (which can be a premium in certain cases). These discounts or premiums depend on the crude oil quality (API Gravity, sulphur content, etc) and can vary through time. The price discounts used in the model are the long term discounts forecasted by Rystad and the reference oil price is the Brent price. Table 3 displays the reserve weighted average discount by oil type. All types are sold on average at a discount as compared to the Brent price. This can be explained by the high quality of the oils constituting the Brent crude stream. The average 12% discount across all reserves implies that while using a reference oil price of $80/bbl, selling prices are on average closer to $70/bbl in the model. Heavier oils being more difficult to refine in order to yield the same products, their discounts are relatively higher. NGL is a particular case, being sold at approximately half the price of oil.

2.5 Counterfactual scenarios

Main scenarios. The analysis relies on the comparison between two scenarios: one in which the carbon tax is implemented in 2016 (the "*Baseline*" scenario), and another in which the tax is delayed until 2030 without being anticipated by investors (the "*No anticipation*" scenario). In terms of investment and production decisions, this second scenario is equivalent to a "business-as-usual" until 2030, when the regime suddenly changes and the tax is implemented.

Alternative anticipations by investors. Additional policy delay scenarios are assessed in this study. These scenarios are designed to span a plausible range of investor beliefs about the timing and stringency of future carbon pricing. In the "*Tax in 2030*" scenario, investors anticipate the correct level of tax to be implemented in 2030, and thus anticipate the social planner behaviour correctly. In the "*5 years after startup*" scenario, investors anticipate the correct level of tax to be implemented after 5 years of production in the asset. This is one of the methods used in the oil and gas sector to screen projects before investing in it. In the "*75% of the tax*" and "*50% of the tax*" scenarios, investors anticipate the tax to be implemented directly after investing in an asset, thus being imposed on the totality of its lifetime production, but underestimate its level by 25% and 50%. This approach seems to closely match reality: Ahluwalia (2017) shows that most oil and gas companies verify that their projects could sustain the introduction of a tax, but usually use

tax rates far lower than the ones needed to be in line with the 1.5°C or 2°C objectives. The 50% underestimation appears close to observed investors' behaviour.

Heterogeneous firms' beliefs Finally, four scenarios are added, in order to assess the effect of potential heterogeneity across firms in terms of anticipations. For two specific subsets of assets, investments are assumed to be screened assuming an immediate implementation of the full tax (thus matching the "Baseline" scenario), or an implementation of the tax in 2030 (similar to the Tax in 2030 scenario). The first subset of assets correspond to oil projects in which European firms own at least 10% of the shares. This is justified by the relatively higher pressure put on firms based in European countries in terms of environmental responsibility. The second subset is all assets owned at 10% by a firm publicly listed. Publicly listed companies can face higher pressure by investors to display greener behaviour and prove resilient to carbon taxation.

2.6 Key metrics

Carbon budget overshoot. The first objective of the paper is to estimate the carbon budget overshoot stemming from the delay in the implementation of the carbon tax. The overshoot is computed by comparing cumulative emissions in one of the delay scenarios with the one that would be observed in the "Baseline" scenario. Once the overshoot is computed for each scenario and level of tax, I disentangle between both the intensive and the extensive margins. The approach there is straightforward: overproduction is considered to happen at the intensive margin if the corresponding asset produces any oil in the "*Baseline*" scenario or produced any oil before 2016. It is considered as the extensive margin if the asset does not produce any oil in the "*Baseline*" scenario. Since the set of anticipations only affect investment decisions, the results found with the different scenarios only vary at the extensive margin.

A reference level of tax used throughout the paper is the tax necessary to match the carbon budget for oil set by Welsby et al. (2021). To adapt it to the framework of this paper, the 747Gbbl budget for oil over the 2018-2100 period is adapted to the 2016-2100 period by adding observed production in 2016 and 2017. It is then converted into a carbon budget by multiplying it by the average downstream oil carbon intensity used in most models (including the TIAM-UCL used in Welsby et al. (2021)), derived from IPCC (2006). The

resulting carbon budget for oil in 2016 is equal to 350.3 GtCO2eq. A shortcoming of this study is that one could legitimately expect the carbon budget dedicated to oil to be reduced should carbon taxation be imposed on oil, and leak towards other fossil fuels. Results are thus to be interpreted as if similar policies were applied to other fossil fuels production in parallel. In order to assess the potential impact of this assumption on my results, I include in robustness checks scenarios in which the carbon budget is reduced by 10% and 25%.

Infrastructure lock-in. Assessing lock-in effects requires to switch to a commitment accounting perspective. The objective is to identify for each asset the year of investment that locks it in into full development. For a given asset and set of anticipations, the year of investment responsible for the locking-in of the asset can vary depending on the level of tax. In order to identify the investments responsible for the lock-in of emissions, I compute for each asset the maximum amount of remaining development costs $\hat{C}_{iT_0}(\tau)$ for the asset to be irreversibly developed. Rearranging equation 1:

$$\bar{C}_{iT_0}(\tau) \;=\; max_{\{T\}} \sum_{t=T_1}^{T} \frac{P(1-d_i)q_{it} \;- c_i}{(1+\delta)^{t-T_0}} \;- \sum_{t=T_\tau}^{T} \tau \;\theta_i \; q_{it} \frac{(1+r)^{t-2016}}{(1+\delta)^{t-T_0}} \tag{3}$$

The year of investment responsible for the lock-in of the asset is the year during which this threshold is crossed, i.e. the investment that makes remaining development expenses lower than this threshold. This year of investments "responsible" for the lock-in means that even if the tax was to be implemented immediately after, the asset would still be developed all the way. Note that since $\bar{C}_{iT_0}(\tau)$ depends on τ, the year of investments "responsible" for the asset development can vary depending on the level of tax, meaning that for different levels of tax, the overcommitted emissions from one asset can be attributed to different years. This also means that investments made during the delay can lock-in assets which production will start after the tax implementation, sometimes multiple years later.

(Over)committed emissions. Committed emissions refer to the cumulative future emissions that are expected to occur from an asset over its remaining lifetime once it has become locked-in. The quantity of committed emissions depends on the level of carbon tax, since it determines the timing of phasing-out of the asset. Overcommitted emissions are the subset of committed emissions attributable to lock-in effects induced by the policy delay: they correspond to the committed emissions from assets whose development would have

been deterred under timely policy implementation, but which, once the lock-in threshold is crossed during the delay, will be developed and produce even if the tax is implemented immediately afterwards. These overcommitted emissions can span over multiple decades if the "depth" of the lock-in, i.e. the gap between the tax necessary to deter the development of the asset and the tax required to phase out production once it has started, is large.

Tax lock-in premium and "naïve" overshoot. I define the naïve overshoot as the residual carbon budget overshoot that arises when the delayed tax is calibrated using the pre-delay supply curve, ignoring the reduction in breakeven costs induced by new investments during the delay. The carbon tax lock-in premium is the wedge between this naïvely calibrated tax and the higher tax needed in the implementation year to offset the full extent of that overshoot.

3 Results

3.1 Overshoot of the budget due to the delay.

I first examine the cumulative carbon budget overshoot (relative to the 2016 benchmark) for a given level of carbon tax delayed from 2016 to 2030. Figure 1 depicts the corresponding carbon budget for each level of tax implemented in 2016 (baseline emissions), as well as the overshoot resulting from the delay at both the intensive and the extensive margins. The total overshoot increases slowly with the tax up to \$50/tCO2eq, before increasing steadily at a faster rate once the tax reaches higher levels. At low tax levels, cumulative production and in turn carbon budget overshoots are only marginally affected, highlighting the high levels of rent perceived by oil producers.

The magnitude of the overshoot at each margin does not evolve in the same way with the carbon tax. As expected, it increases at the intensive margin as the tax gets higher. This is a mechanical effect due to the fact that the higher the tax, the larger the share of production between 2016 and 2030 that would be driven out of the market in the presence of the tax. However, the effect at the extensive margin does not display such a simple behaviour. It increases steadily as the tax goes higher, and reaches a plateau once the tax exceeds \$100/tCO2eq, near its required level in order to respect the 1.5°C carbon budget.

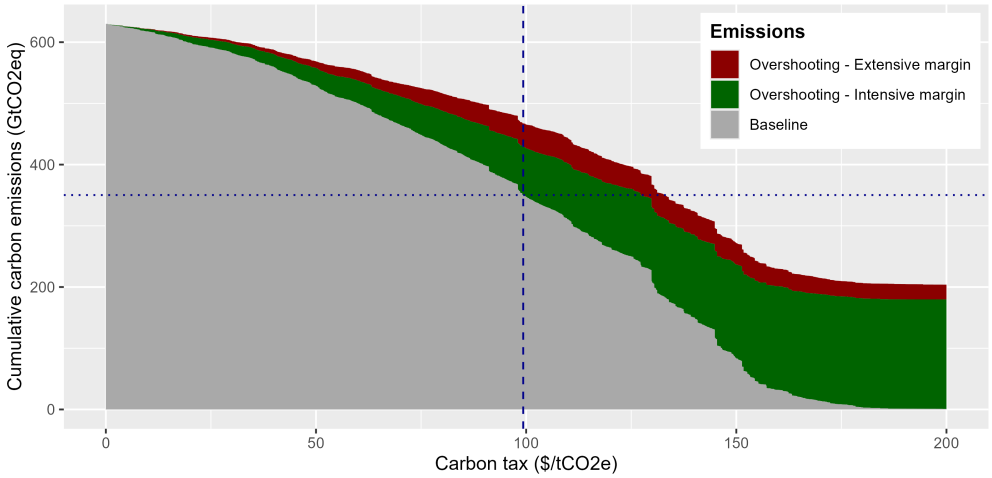

Figure 1: Oil carbon budget overshooting when delaying carbon taxation from 2016 to 2030.

Notes: The x-axis reports levels of carbon tax. The grey area corresponds to the ""Baseline" scenario" scenario, i.e. the cumulative emissions if the tax is implemented in 2016. The green and red areas depict the overshoot resulting from delaying the tax until 2030 at the intensive and extensive margins respectively. The dotted horizontal line indicates the oil carbon budget set by Welsby et al. (2021). The dashed vertical line is the corresponding carbon tax needed in 2016 in order to constrain cumulative emissions within this budget. The scenario used to compute the overshoots is the "no anticipation" one. All computations are made assuming a reference Brent oil price of $80/bbl.

At these levels, overemissions at the extensive margin fluctuate around 40 GtCO2eq. This feature can be explained by the relatively flat nature of the supply curve around the cheapest oil sources. For an asset to be included in the overshoot at the extensive margin, it must be that its project breakeven price is too high to enter the market in presence of the tax, but its production breakeven price is low enough so as to start production if it is developed. In addition to high-capital-intensity assets, will also be included assets whose project breakeven prices are just above the threshold to enter the market. These assets are thus more numerous when the threshold is situated around the levels of price where the supply curve is flatter, i.e. for levels of tax driving out a significant share of the market. Any ambitious carbon taxation delayed would thus lead to a significant overshoot at the extensive margin.

The dashed and dotted blue lines in figure 1 depict the oil carbon budget set by Welsby et al. (2021) updated as of 2016 (350.3GtCO2eq), and the corresponding carbon tax necessary in 2016 in order to constrain cumulative emissions within it ($99.3/tCO2eq). For this level of tax, the magnitude of the overshoot amounts to 78.7 GtCO2eq at the intensive margin and 38.4 GtCO2 at the extensive one, which combined represent 33% of the initially set carbon budget. To provide a better perspective, it is worth noting that a 117 GtCO2eq overshoot is equivalent to nearly three years' worth of global emissions as estimated in 2020 (Climate Watch, 2023). $99.3/tCO2eq (hereafter approximated to $100/tCO2eq in the paper) is the

level of tax used as reference in the remaining of this paper when assessing the timing of overproduction and the effect of carbon lock-in.

Supply source heterogeneity. I further decompose in figure 2 the overshoot by oil supply segment (main types of oil sources) in order to determine which types of assets contribute the most to each margin. First, it is worth noticing that in the ""Baseline" scenario" scenario, a $100 tax drives out of the market most unconventional productions: oil sands and shale oil represent together 0.3% of total cumulative supply. The high operating costs of oil sands and shale assets entail that most of their overproduction happens at the intensive margin. The higher carbon intensity of oil sands further reinforces this effect. More generally, the overshoot at the intensive margin is heavily concentrated in onshore assets. On the opposite end, deepwater assets have high development costs but are relatively cheap to operate, meaning that their overproduction is largely due to additional assets being developed during the delay. Similar conclusions can be drawn regarding midwater assets, to a lesser extent. This translates into offshore assets being responsible for two third of the overshoot at the extensive margin, with around 60% of it from deepwater assets alone.

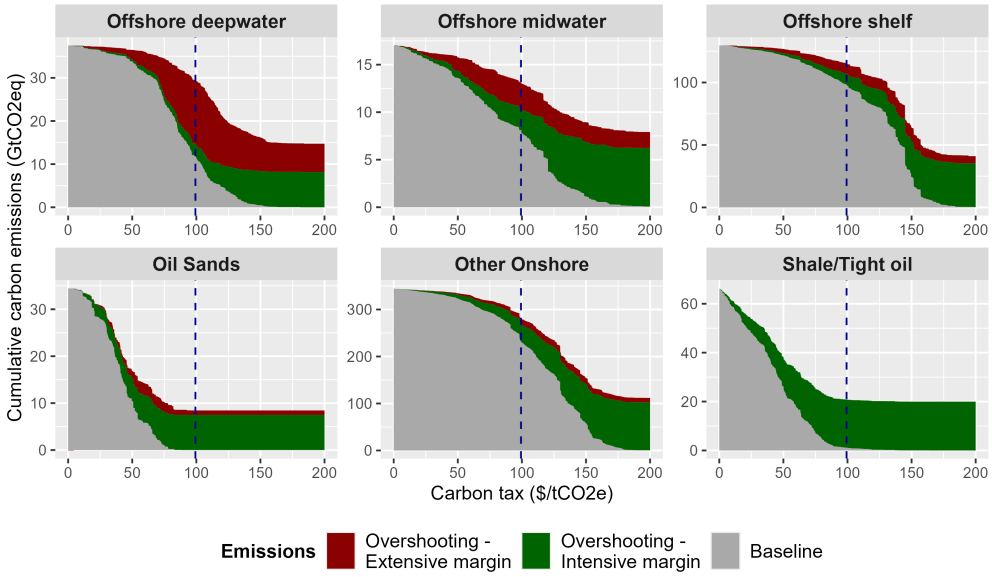

Figure 2: Oil carbon budget overshooting by supply segment.

Notes: Each panel corresponds to an oil supply segment. For each panel, the x-axis is the level of carbon tax. The grey area corresponds to the ""Baseline" scenario" scenario, i.e. the cumulative emissions if the tax is implemented in 2016. The green and red areas depict the overshoot resulting from delaying the tax until 2030 at the intensive and extensive margins respectively. The dotted horizontal line indicates the oil carbon budget set by Welsby et al. (2021). The dashed vertical line is the corresponding carbon tax needed in 2016 in order to constrain cumulative emissions within this budget. The scenario used to compute the overshoots is the "no anticipation" one. All computations are made assuming an oil price of $80/bbl.

Regional heterogeneity. This heterogeneity in terms of oil sources has implications for the geographical distribution of the overshoot. Figure 19 displays its decomposition by world region. Comparing North and South America provides a striking illustration of the regional heterogeneity in terms of overshoot. Because North America concentrates most shale and oil sands production, it alone is responsible for half of the overshoot at the intensive margin. Conversely, a third of the overshoot at the extensive margin is attributable to South America alone, due to the numerous and massive deepwater projects currently being developed along the coasts of Brazil and Guyana. In other regions, each margin participates in similar proportions to the overall overshoot, although its magnitude vary greatly across regions. Middle Eastern production is not heavily affected by either the tax or the delay in implementation, while North and South American production greatly depends on the delay and would virtually disappear in presence of the tax. Asian, European and Russian productions are also greatly affected by the tax. This pattern may partly explain the focus put by Occidental countries on demand-side policies rather than supply-side ones.

Overshoot at the extensive margin by scenario of anticipations. Ahluwalia (2017) showed that most oil and gas producers anticipate some taxation to be implemented at some point, despite the levels of tax anticipated being low. I investigate the impact of a variety of anticipations scenarios on the overshoot at the extensive margin. I focus on this sole margin because anticipations only play a role in the decision to invest or not in a new project, thus only impacting the extensive margin. Figure 3 displays the overshoot resulting from each scenario aggregated by 5-years time period. A first observation is that anticipations of future tax by producers do not lead to a massive decrease in the magnitude of the overshoot.

The first two alternative scenarios investigate the effect of anticipating the correct level of tax to be implemented later in time. When investors anticipate it to be implemented in 2030 (thus perfectly predicting the delay) or 5 years after production starts, the overshoot at the extensive margin still represents around 70% of the one in the "no anticipation" scenario. While not greatly affecting the total magnitude of the overshoot, the two sets of anticipations lead to a different time distribution: anticipating a tax in 2030 phase out most investments in new assets in the 2025-2029 period, thus leading to a smaller overproduction after the end of the delay. Conversely, the effect of anticipating the tax to be implemented 5 years after

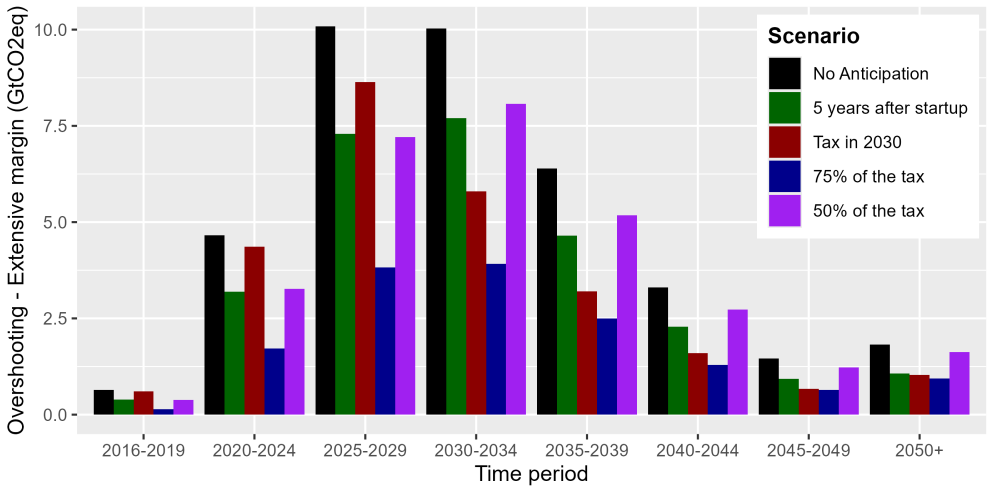

Figure 3: Yearly oil carbon budget overshooting at the extensive margin by scenario.

Notes: 5-years periods are on the x-axis. "2050+" corresponds to the 2050-2099 period. The y-axis corresponds to the quantities of oil overproduced by period and by scenario of anticipation compared with the "Baseline" scenario . The level of tax used in the graph is \$99.3/tCO2eq. All computations are made assuming a reference Brent oil price of \$80/bbl.

production startup affects investment decisions uniformly during the whole delay, leading to a reduction of the overshoot more or less constant over time.

The two following scenarios assume that investors expect the tax to be implemented any time, but underestimate its level. I provide results for taxes equivalent to 50% and 75% of the optimal one. Underestimating the tax by only 25% reduces overemissions at the extensive margin significantly (60% decrease compared with the no anticipation scenario), although the decrease in not proportional to the tax relative to its optimal level. This is confirmed when assuming a 50% underestimation of the tax. In this case, less than 23% of overemissions at the extensive margin disappear. While these two last scenarios replicate the behaviour of investors as observed by Ahluwalia (2017), the 25% underestimation is probably the least realistic one. Only few producers screen the viability of their projects in the presence of a tax higher than \$60/tCO2eq, and the sole purpose of these high tax screening is to stress-test existing assets, not guide investments in new assets. Most carbon shadow prices used in the oil industry to screen new projects lie in the \$23-\$60 range. The 50% underestimation scenario is therefore much more aligned with investors anticipations described by Ahluwalia (2017). This explains the relative continuity of oil investments observed despite more and more firms adopting internal carbon pricing schemes.

Heterogeneous beliefs across firms. Figure 4 reports the same results considering

various scenarios of heterogeneous behaviour across firms. I test a first scenario in which only a subset of investors (European firms or public firms) screen their investments with the full tax, as in the "Baseline" scenario, and another one in which the same subset of investors anticipate the full tax to be implemented in 2030, as in the "Tax in 2030" scenario. For reference, approximately 20% of assets and reserves are at least partially owned by European firms, while the share of assets with publicly-listed firms shareholders represents two thirds of the total. Unsurprisingly, screening projects with a $100 tax reduces greatly the overshoot at the extensive margin, by 40% for the Europe-owned subset and a striking 84% for the public subset. However, in line with results shown in figure 3, anticipating the tax to be put in place in 2030 only slightly affects its magnitude.

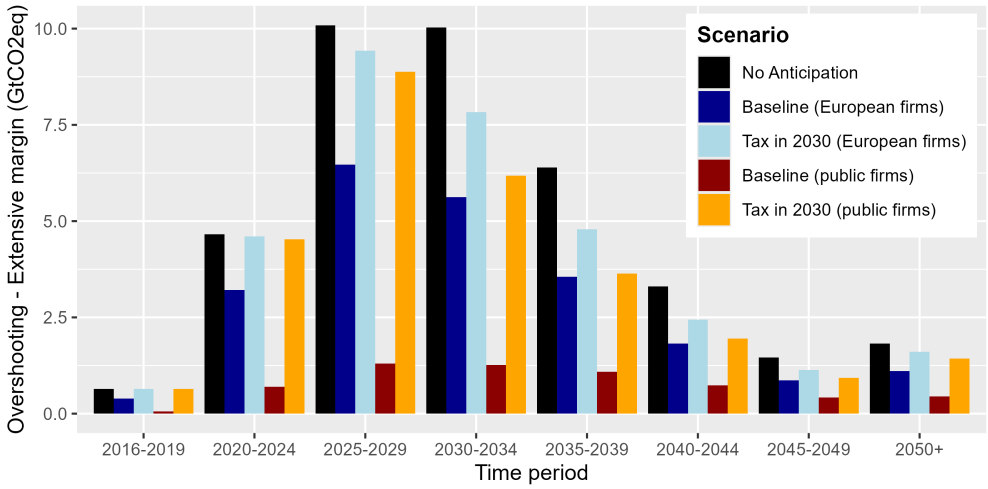

Figure 4: Yearly oil carbon budget overshooting at the extensive margin by scenario of firms heterogeneous beliefs.

Notes: 5-years periods are on the x-axis. "2050+" corresponds to the 2050-2099 period. The y-axis corresponds to the quantities of oil overproduced by period and by scenario compared with the baseline. The level of tax used in the graph is $99.3/tCO2eq. All computations are made assuming a reference Brent oil price of $80/bbl.

Over-emissions post-2030. A key result is that overproduction remains substantial after the delay: depending on the scenario, 47% to 64% of the overshoot at the extensive margin occurs after the tax is implemented, with 13% to 27% still remaining in 2040, 10 years after its implementation. Focusing on overproduction until the right policy is put in place when studying the delay does not allow to capture this future overproduction, thus leading to an underestimation of the full extent of its cost. The accumulation of investments in new oil assets should not be viewed solely as an increase in production capacity, but rather as committed emissions from infrastructures costly to phase out once they are built.

Next subsection approaches overproduction at the extensive margin through this concurrent perspective, in order to better understand the consequences of climate inaction in terms of overinvestment rather than overproduction.

3.2 The role of infrastructure lock-in and the carbon tax premium

In this section, I shift my analysis from examining the annual emissions associated with each scenario to focusing on carbon lock-in: the cumulative emissions committed by investments in new oil infrastructure during the delay. Unlike annual overshoot, lock-in quantifies future emissions that become unavoidable once assets are developed, even if policies tighten later. This approach allows for a more comprehensive assessment as it quantifies the future emissions attributed to each year of delay. To evaluate the quantities of committed emissions, I define the lock-in threshold as the cumulative investment beyond which an asset's development becomes irreversible under a given tax level. The year in which this threshold is reached is considered to be the year of investment responsible for the lock-in of this asset. The analysis reveals that more than half (51%) of all committed emissions during the delay are overcommitted (i.e. attributable to the delay itself) and that these emissions are disproportionately concentrated in a small number of mainly offshore projects. This highlights how few capital-intensive projects drive most carbon lock-in effects, with major implications for targeted climate policies.

Figure 5 reports committed emissions by supply source, split between committed emissions existing in the "Baseline" scenario and overcommitted emissions. As expected, virtually no new unconventional project is developed during the delay. 70% of overcommitted emissions are concentrated in offshore assets, with deepwater assets representing the majority of it. More than 80% of committed emissions in these projects are overcommitted ones, compared with around 40% for onshore and offshore shelf projects.

The depth of lock-in. Reported overcommitted emissions in figure 5 are relative to a single tax level of $100/tCO2eq. However, it does not give information on the depth of the lock-in. If assets responsible for overcommitted emissions only required a slightly higher tax in order to be phased out, increasing the tax to compensate for it would not entail a redistribution of stranded assets across the globe. However, if these assets are resilient to high taxes, the necessary tax increase to compensate for the delay would involve stranding

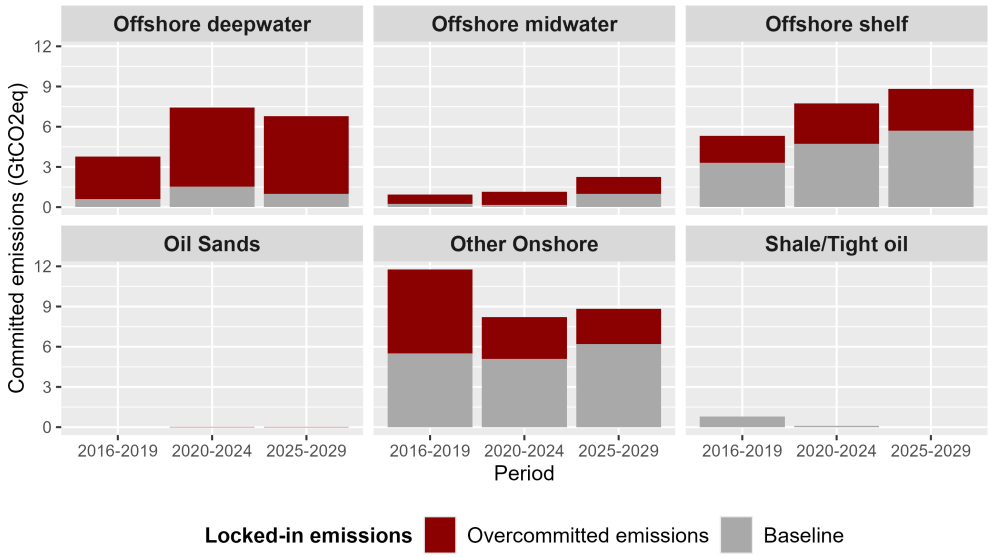

Figure 5: Overcommitted emissions by oil supply segment.

Notes: Each panel corresponds to an oil supply segment. Each bar corresponds to the amount of committed emissions from investments in the development of new oil assets in the corresponding segment.The grey share of the bar depicts committed emissions that are emitted in the "Baseline" scenario. The red share corresponds to committed emissions from assets that are not developed in the "Baseline" scenario, thus representing the cost arising from carbon lock-in effects. Committed emissions from a given asset are attributed to the year from which the full development of this asset becomes inevitable, which depends on the level of tax. The level of tax used for this graph is \$100/tCO2eq. The scenario used for this graph is the "No anticipation" one. All computations are made assuming a reference Brent oil price of \$80/bbl.

assets that would not have been affected by the tax should it be implemented in 2016. Figure 6 maps projects responsible for overcommitted emissions, aggregated by province. The color of the dots reports the average level of required tax to phase out assets in the province. While a small number of assets, mainly in China, can be phased-out with taxes below \$120/tCO2eq, most overcommitted emissions would remain even at high tax levels (more than 86% at \$120/tCO2eq and still 41% at \$140/tCO2eq). Delaying climate policies is thus not neutral in terms of geographical distribution of stranded assets.

Lock-in premium and naïve overshoot. The extensive overproduction occurring during the delay entails sharp year-to-year increases in the optimal carbon tax needed to meet the 1.5°C objective. Beyond this direct overproduction effect, cumulative investments in new production facilities lead to a downward shift in the supply curve and higher locked-in emissions for any given level of tax. Accounting for investment dynamics therefore further increases the tax required to constrain emissions within the carbon budget. I call this gap between the "naïve" carbon tax and the optimal one the *carbon tax lock-in premium*. It is defined as the difference between (i) the carbon tax that would be sufficient to meet the

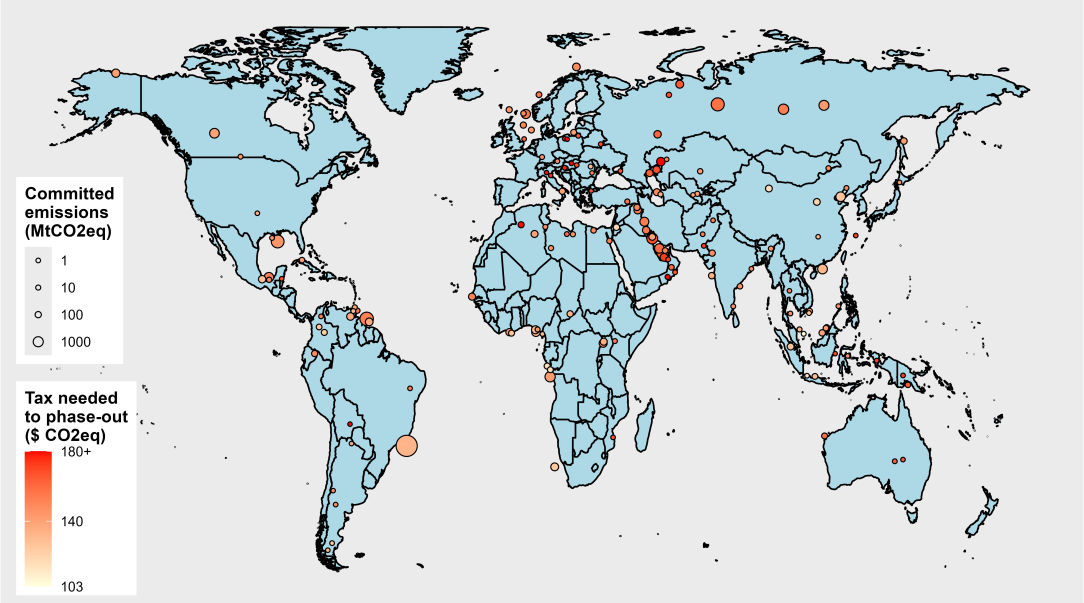

Figure 6: Carbon tax required to phase-out projects responsible for overcommitted emissions (aggregated by province).

Notes: Each circle correspond to one province in which overcommitted emissions exist. A province is a level of disaggregation available in Rystad's data situated between the asset-level and the country-level. Circle sizes depend on the quantities overcommitted and colors reflect the average tax needed to phase-out committed emissions right after commitment. The scenario used for this graph is the "No anticipation" one. All computations are made assuming a reference Brent oil price of $80/bbl.

carbon budget if this drop in asset breakeven prices was ignored, and (ii) the higher tax required once investment-driven lock-in effects are accounted for.

The lock-in premium depends on the initial state of oil production facilities, and builds up as new projects are developed (and their breakeven price decreases). Ignoring the change in the supply curve induced by these new investments, the naïve tax needed in 2030 in order to remain within the carbon budget is set at $130/tCO2eq. However, implementing this tax from 2030 onwards would still lead to overshooting the budget by more than 18 GtCO2eq (12.5% of the remaining budget in 2030), due to overcommitted emissions from assets developed during the delay. Resolving this residual overshoot caused by infrastructure lock-in requires adding a lock-in premium of $3/tCO2eq to the tax in order to effectively meet the budget, resulting in an optimal 2030 carbon tax of $133/tCO2eq[10]. Ignoring infrastructure lock-in thus leads policymakers to systematically underprice carbon, even when the carbon budget itself is correctly identified.

[10]Note that implementing such tax in 2016 would lead to cumulative emissions equal to 183 GtCO2eq, roughly 48% below the 1.5°C carbon budget. This underscores the strikingly high cost of delaying carbon taxation, even for a 14 years period.

Table 1 shows year-to-year change in the carbon tax required to meet the carbon budget, increasing from \$99.3/tCO2eq in 2016 to \$133/tCO2eq in 2030. The carbon tax lock-in premium is indicated for each year, as well as the overshooting of the budget resulting from the implementation of the naïve tax (the percentage between parenthesis is relative to the remaining carbon budget). Results are provided for two periods of delay and investments: the "Baseline" scenario 2016-2030, as well as the 2025-2030. The initial state of existing facilities and production costs in the second period is thus the one observed in 2025.

As outlined in the table, the magnitude of the lock-in premium does not necessarily reflect the real climate cost of implementing the naïve tax. Its maximum level of \$6.4 is reached in 2026 whereas overemissions associated with the naïve tax rank higher in 2030. This can be explained by the fact that the supply curve can be locally flat, with large reserves becoming locked-in at similar tax rates[11].

			From 2016		From 2025	
Tax year	Remaining budget ($GtCO_2eq$)	Lock-in adjusted carbon tax ($/$tCO_2eq$)	Tax premium ($/$tCO_2eq$)	Naïve overshoot ($GtCO_2eq$)	Tax premium ($/$tCO_2eq$)	Naïve overshoot ($GtCO_2eq$)
2016	350.3	99.3	–	–	–	–
2018	321	105	2.6	7.3 (2.3%)	–	–
2020	291.8	109.7	2.4	9.6 (3.3%)	–	–
2022	264.6	112.7	2.1	12.4 (4.7%)	–	–
2024	236	118	3.5	13.3 (5.6%)	–	–
2026	206.9	125.5	6.4	16.4 (7.9%)	0.6	2.7 (1.3%)
2028	177	129.8	2.6	15.5 (8.7%)	<0.1	7.3 (4.1%)
2030	146.7	133	3	18.4 (12.5%)	2.2	10.1 (6.9%)

Table 1: Lock-in premium and naïve overshoot by tax implementation year (constant tax).

Notes: The table reports the remaining carbon budget and the optimal carbon tax required to respect it for each implementation year. The tax premium corresponds to the additional carbon tax required due to investment-driven lock-in. Overshoot reports the implied carbon budget overshoot in $GtCO_2eq$, with percentages expressed relative to the remaining budget. "From 2016" and "From 2025" indicate the baseline year from which delay is measured.

The buildup of lock-in. The cost of not accounting for infrastructure lock-in increases at rapid pace, especially over the first few years. The relatively lower buildup af-

[11] A striking example of this is shown in the case of the 2025-2030 period. The required tax premium in 2028 is close to 0 while being associated with a 7.3GtCO2eq naïve overshoot. This is generally caused by a unique asset with particularly large locked-in reserves.

terwards is due to the fact that additional overcommitted emissions from new investments are partly compensated by past overcommitted emissions being effectively emitted. This explains the fast increase in the naïve overshoot when starting from 2025. Overlooking only five years of investments would entail exceeding the budget by more than 10 GtCO2eq, almost 7% of the remaining budget.

Crucially, the overcommitted emissions constituting this remaining overshoot are defined under a much more restrictive counterfactual than the ones presented in the former section. In both cases, overcommitted emissions originate from assets whose development is dependent on the \$100/tCO2eq tax being delayed to 2030. However, while committed emissions in the former section corresponded to locked-in emissions in a \$100 tax world, the overshoot computed in table 1 corresponds to locked-in emissions in a world where the much higher naïve tax is implemented. It is a direct consequence of the deep gap between full project breakeven prices and production breakeven prices outlined in figure 6.

3.3 Increasing tax

This section presents the main results outlined in the precedent section, this time considering an increasing carbon tax. Under the constraint of a carbon budget, the dynamically efficient carbon tax increases at the discount rate (Lemoine and Rudik, 2017; Dietz and Venmans, 2019). Additionally, the marginal damage of emissions increases over time (Greenstone et al., 2013). A third argument in favour of such tax is its higher acceptability by the public and the industry, giving it time to adjust. The tax assessed in this section increases at a 2% social discount rate[12]. The choice of a discount rate lower than market rates is justified by uncertainty about the far-distant future growth and interest rates (Stern, 2007; Weitzman, 2007). The declining discount rate literature points to social discount rates ranging from 1% to 2% (Weitzman, 1998; Gollier, 2009), while Nordhaus (2007) justifies using a higher rate, around 3%. I adopt a social discount rate of 2% which lies between these different approaches. Different rates ranging from 1% to 6% are included in the sensitivity analysis.

Overshoot due to the delay. Figure 7 reports the overshoot at both the intensive and extensive margins by initial carbon tax level. The carbon tax consistent with the budget

[12]All tax levels in this section are expressed in initial 2016 level, regardless of the timing of implementing.

from Welsby et al. (2021) is implemented in 2016 at 57.1$/tCO2eq[13]. For this level of tax, the delay entails a significant overshoot of the budget, though smaller than the one computed with a constant tax, amounting to 65 GtCO2eq GtCO2eq in total (18.6% of the budget), of which 36% happens at the extensive margin. This smaller amount can be explained by the relatively low level of tax during the delay, reducing mechanically the overshoot at the intensive margin. The extensive margin is also reduced because most assets developed during the delay would be economically viable in a low tax environment, with the high rentability of oil projects allowing to compensate for development costs in only few years of production. Assets developed during the delay are also phased-out earlier than in a constant tax environment, when the tax gets too high for operations to continue, further reducing the post-2030 overshoot at the extensive margin. Yet, delaying would still entail exceeding the carbon budget by a considerable margin, and allow for the development of many unnecessary oil projects.

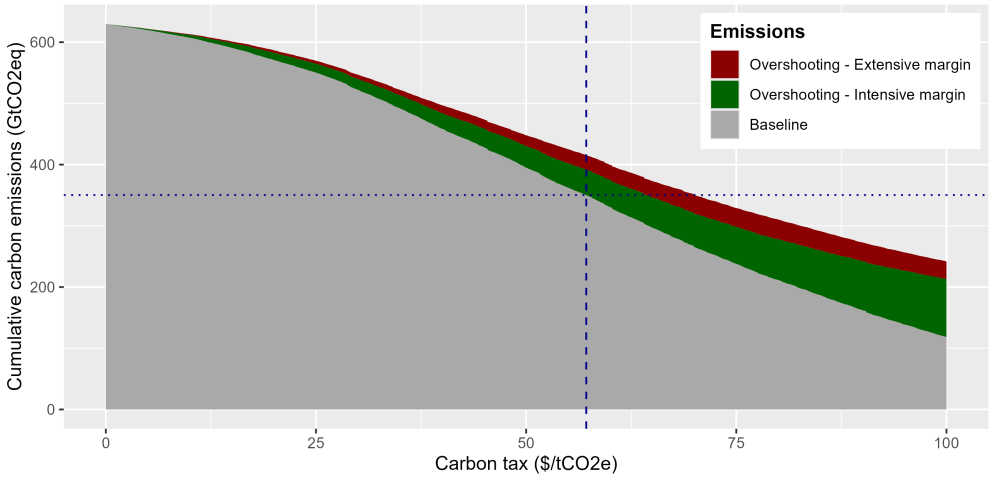

Figure 7: Oil carbon budget overshooting from the delay - 2% increasing tax.

Notes: The x-axis is the 2016 level of carbon tax, increasing annually at a 2% rate. The grey area corresponds to the ""Baseline" scenario" scenario, i.e. the cumulative emissions if the tax is implemented in 2016. The green and red areas depict the overshoot resulting from delaying the tax until 2030 at the intensive and extensive margins respectively. The dotted horizontal line indicates the oil carbon budget set by Welsby et al. (2021). The dashed vertical line is the corresponding carbon tax needed in 2016 in order to constrain cumulative emissions within this budget. The scenario used to compute the overshoots is the "no anticipation" one. All computations are made assuming a reference Brent oil price of $80/bbl.

Heterogeneity across sources. Figures 21 and 21 report the overshoot by supply segment and regions respectively. Most conclusions drawn in the case of a constant tax stay valid when assessing an increasing one. Oil sands and shale concentrate a large share of

[13]For reference, the tax amounts to $75/tCO2eq in 2030, reaches $100/tCO2eq in 2045 and $200/tCO2eq in 2080.

overproduction at the intensive margin, while deep offshore projects remain a major source of overshoot at the extensive one. North and South America still play complementary roles in explaining the total overshoot, with North American assets overproducing during the delay and South American ones after.

A notable difference is the relatively smoother reduction in emissions induced by increasing the level of tax. All oil and gas producers are affected more uniformly, mainly in later years of the period when the tax reaches high levels. This translates into a higher reduction in cumulative emissions from regions only marginally affected by the constant tax, such as Middle Eastern countries. Yet, Middle Eastern oil projects do not appear to give rise to overproduction at any margin. A probable explanation is their low production cost and carbon intensity. Most Middle Eastern assets are resilient to tax levels during the delay and only get affected by the policy in later years.

Lock-in premium and naïve overshoot. Table 2 reports the carbon tax lock-in premium by year as well as the overshoot associated with implementing the naïve tax. Ignoring investment-driven lock-in from 2016 results in a residual overshoot of the budget of 17.1 GtCO2eq, comparable to the one found with a constant \$100/tCO2eq tax. The tax premium required to offset it amounts to 3.7\$/tCO2eq, expressed in 2016 levels (5.1\$/tCO2eq in 2030). Ignoring the next five years of investments leads to a similar cost to the one computed for a constant tax, around 10 GtCOeq. Overall, although the tax increase rate affects the distribution of stranded assets, the magnitude of the cost resulting from ignoring lock-in effects remains unchanged.

Despite reducing the lifetime of assets, in turn reducing the quantity of committed emissions of newly developed assets, the naïve overshoots remains consequent when considering a 2% increasing tax. A consequence of the increasing nature of the tax is the fact that almost all assets responsible for overcommitted emissions will have a share of cumulative emissions locked-in once the tax is in place. The tax being relatively low in the 2030-2039 decade, most locked-in assets will continue producing, and increasing marginally the level of tax can only induce slightly earlier phase-out. Consequently, this requires to assess the depth of lock-in using a different and more restrictive approach compared with the one presented in figure 6.

The map in figure 8 displays the size of overcommitted emissions from assets undevel-

Tax year	Remaining budget ($GtCO_2eq$)	Lock-in adjusted carbon tax ($/$tCO_2eq$)	From 2016		From 2025	
			Tax premium ($/$tCO_2eq$)	Naïve overshoot ($GtCO_2eq$)	Tax premium ($/$tCO_2eq$)	Naïve overshoot ($GtCO_2eq$)
2016	350.3	57.1	–	–	–	–
2018	321	58.3	0.4	2.8 (0.9%)	–	–
2020	291.8	59.3	0.7	5.1 (1.7%)	–	–
2022	264.6	60.4	0.9	8.2 (3.1%)	–	–
2024	236	62.6	2	10.9 (4.6%)	–	–
2026	206.9	64.6	2.5	13.8 (6.7%)	0.4	1.9 (0.9%)
2028	177	67.2	3.2	15.1 (8.5%)	1.3	5.6 (3.2%)
2030	146.7	70	3.7	17.1 (11.6%)	2	10.4 (7.1%)

Table 2: Lock-in premium and naïve overshoot by tax implementation year (increasing tax).

Notes: The table reports the remaining carbon budget and the optimal carbon tax required to respect it for each implementation year. Carbon taxes are reported as their 2016 level and increase annually at a 2% rate. The tax premium corresponds to the additional carbon tax required due to investment-driven lock-in, again in 2016 level. Overshoot reports the implied carbon budget overshoot in $GtCO_2eq$, with percentages expressed relative to the remaining budget. "From 2016" and "From 2025" indicate the baseline year from which delay is measured.

oped in the ""Baseline" scenario" scenario, aggregated by province, as well as a measure of this depth of lock-in. Points colour denotes the 2016 level of carbon tax required to reduce by 10% the quantity of overcommitted emissions. The geographical distribution of overcommitted emissions is relatively similar to the case of a constant tax. However, the lock-in appears to be even deeper entrenched in the case of an increasing tax. To the exception of Canadian oil sands and a few relatively marginal onshore projects, the vast majority of locked-in emissions are resilient to carbon taxes higher than 70$/tCO2eq. Implementing a 70$/tCO2eq tax (equivalent to the lock-in adjusted tax necessary in 2030) would leave 82% of overcommitted emissions unchanged, and a 80$/tCO2eq one 67%. Reserves are even more locked-in in deep offshore projects with 86% and 75% of overcommitted emissions resilient to these two levels of tax.

4 Sensitivity

The central quantitative results presented in section 3, a naïve overshoot ranging from 17 to 19 GtCO2eq and a tax lock-in premium of 3 to 4$/tCO2eq, rest on a specific combination of assumptions regarding investor behaviour, cost structure, carbon intensity and

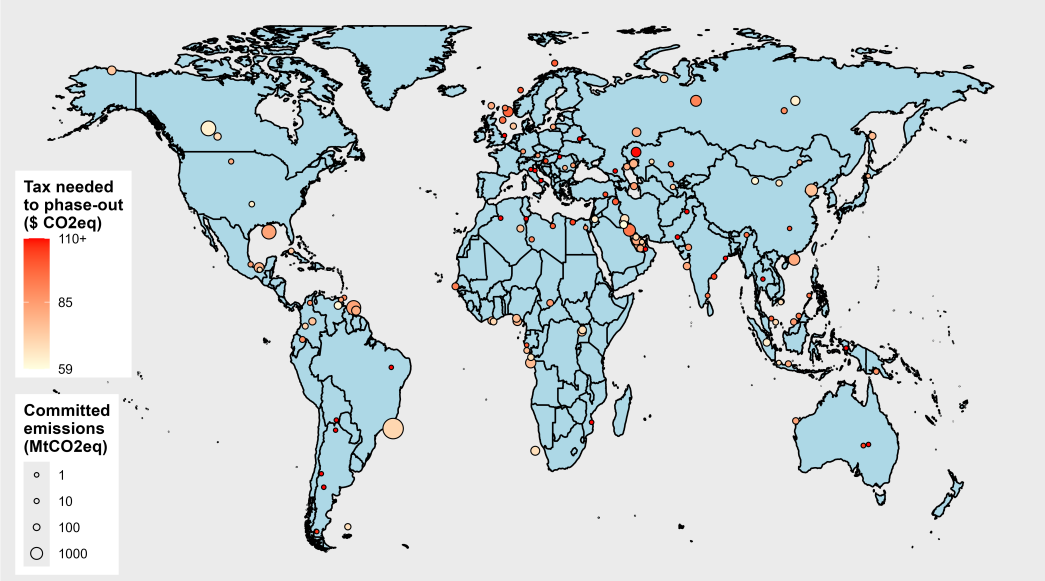

Figure 8: Carbon tax required to reduce overcommitted emissions from locked-in projects by 10% (aggregated by province).

Notes: Each circle correspond to one province in which overcommitted emissions exist. A province is a level of disaggregation available in Rystad's data situated between the asset-level and the country-level. Circle sizes depend on the quantities overcommitted and colors reflect the province average of the 2016 increasing carbon tax required to reduce locked-in emissions in a 57.1$/tCO2eq world by 10%. The scenario used for this graph is the "No anticipation" one. All computations are made assuming a reference Brent oil price of $80/bbl.

carbon tax design. This section evaluates the sensitivity of these results to a variety of assumptions. Results are reported for both the constant and increasing tax cases, and are separated by category of assumptions. The first subsection focuses on the effect of anticipation scenarios, the second on assumptions on costs definition and the computation of NPV, and the third on assumptions on carbon intensities. In the fourth subsection, the impact of various carbon tax increase rate is studied. Finally the last subsection incorporates a time-varying pass-through of the carbon tax between producers and consumers. Across all configurations, the naïve overshoot in 2030 remains significant, and the lock-in premium is positive throughout.

4.1 Scenarios of anticipation

Figure 9 reports the consequences of various investors anticipations scenarios on the 2030 tax premium and naïve overshoot. As expected, investors anticipating a tax to be implemented at some point always reduces the cost of ignoring lock-in effects in both the cases of a constant tax (panel a) as well as an increasing one (panel b). Large heterogeneity

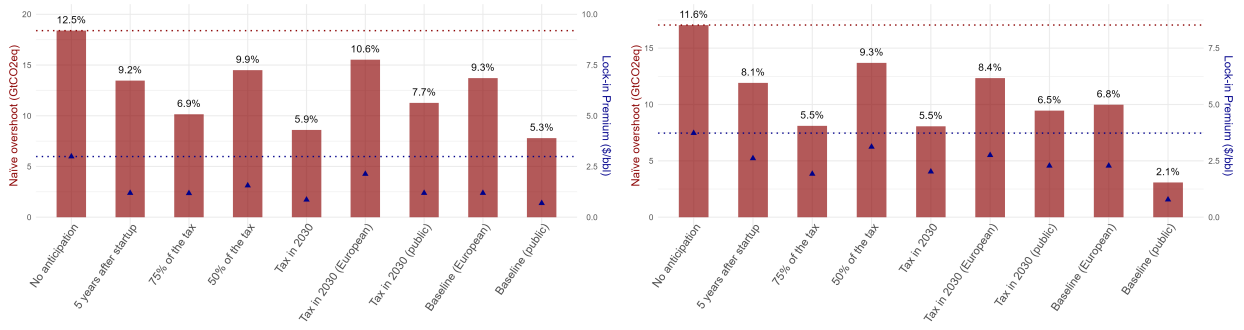

(a) Constant tax

(b) Increasing tax

Figure 9: Sensitivity to scenarios of anticipations.

Notes: Each red bar reflects post-2030 overemissions resulting from the implementation in 2030 of the naïve carbon tax. Above each bar is indicated the corresponding share of the remaining carbon budget in 2030. The red dotted line extends the baseline results. Points correspond to the carbon tax lock-in premium, with the blue dotted line extending the baseline lock-in premium.

exists across scenarios, but the overall results follow the logic highlighted in figures 3 and 4. Screening fully projects but with an underestimated level of tax does drive out of the market a non-negligible share of projects that would otherwise participate in the naïve overshoot of the budget. The scenario closest to the observed behaviour of firms ("50% of the tax") still leads to a sizeable residual naïve overshoot, close to 15 GtCO2eq whatever the tax increase rate. Expecting the tax to be implemented in 2030 dramatically reduces investments realized in later years of the delay, thus reducing significantly the cost of ignoring lock-in effects.

Interestingly, testing heterogeneous beliefs across firms leads to differing conclusions whether the tax increases or not. Although the relative effect of each scenario on the naïve overshoot remains the same, its magnitude is consistently reduced when considering an increasing tax.

4.2 Costs and NPV calculations

I shift my sensitivity analysis to focus on alternative assumptions on development and extraction costs, as well as discount rates. I first test the sensitivity of my results to different selling prices: I use a \$60/bbl and \$100/bbl reference Brent oil price instead of \$80/bbl, and assume in a third configuration that all oils produced are sold at an homogeneous price (thus ignoring asset-specific discounts). I then use a constant production price, defined as the breakeven price of production over the asset full lifetime (still excluding development costs). I then reduce the amount of capital expenses flagged as development ones by excluding facility

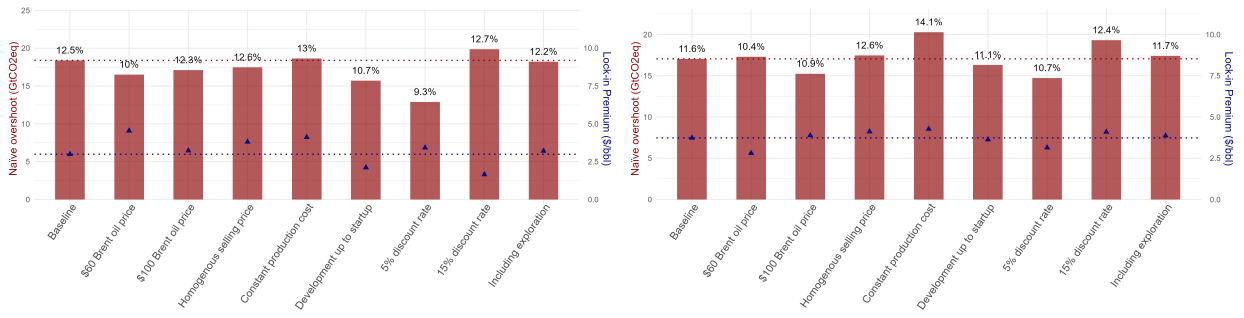

(a) Constant tax

(b) Increasing tax

Figure 10: Sensitivity to assumptions on costs and NPV calculations.

Notes: Each red bar reflects post-2030 overemissions resulting from the implementation in 2030 of the naïve carbon tax. Above each bar is indicated the corresponding share of the remaining carbon budget in 2030. The red dotted line extends the baseline results. Points correspond to the carbon tax lock-in premium, with the blue dotted line extending the baseline lock-in premium.

and well capex over the first two years of production, otherwise included in development expenses. The next two configurations explore the impact of using a 5% and 15% discount rates when computing the net present value of projects. Finally, the last configuration includes all not-yet-discovered assets in 2025, and assimilate their exploration expenses to development ones.

Figure 10 reports the results for each of these configurations in the case of a constant and increasing tax. Overall, the level of the naïve overshoot remains consistent with my baseline results for each of these alternative setups. Unsurprisingly, increasing the discount rate strengthens the gap between breakeven prices including and excluding development expenses, strengthening in turn the lock-in premium and the naïve overshoot. Similarly, reducing the amount of capital expenses included in development expenses mechanically reduces this gap, thus reducing the magnitude of the results. One could expect the inclusion of exploration costs and not-yet-discovered assets in the analysis to increase the total cost of lock-in effects, yet the results do not appear to change significantly. A major reason is that most assets discovered after 2025 will not be developed before the end of the delay, and would therefore be stopped by the introduction of the tax in 2030. They can thus not participate in the buildup of the naïve overshoot.

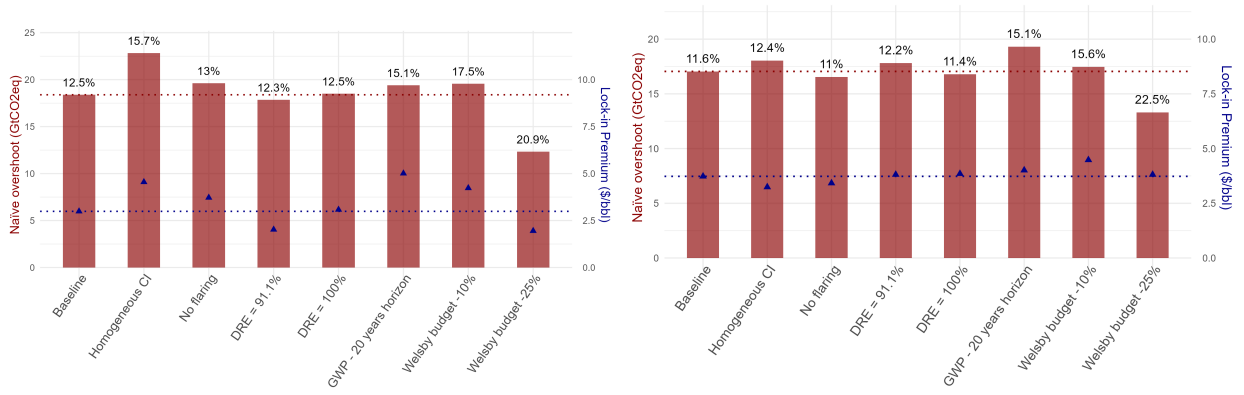

(a) Constant tax

(b) Increasing tax

Figure 11: Sensitivity to assumptions on carbon intensities.

Notes: Each red bar reflects post-2030 overemissions resulting from the implementation in 2030 of the naïve carbon tax. Above each bar is indicated the corresponding share of the remaining carbon budget in 2030. The red dotted line extends the baseline results. Points correspond to the carbon tax lock-in premium, with the blue dotted line extending the baseline lock-in premium.

4.3 Carbon intensities

I further investigate the sensitivity of my results to changing assumptions on the carbon intensity of assets. In the first set of CI, I assume homogeneous CI across assets and assign to all of them a reserve weighted average CI equal to 468 kgCO2eq/bbl. The second set of CI assumes that all assets stop flaring at no cost. The third and fourth sets of CI assume levels of flaring efficiency (DRE, i.e. destruction removal efficiency, see appendix B) of 91.1% and 100% instead of the baseline 96.1%. This greatly affects the quantity of methane being emitted from flaring assets. In the next set of carbon intensity, I use global warming potentials assuming a 20-year horizon instead of the baseline 100-years horizon. This implies sharp increases in CI as well as reordering across assets due to the higher global warming potential of methane associated with this assumption. Finally, I keep the baseline carbon intensity but use a reference carbon budget reduced by 10% and 25%.

Panel 11 reports the results associated with the differing sets of CI and carbon budgets. Results remain in line with my baseline computations for all sets of CI. Interestingly, homogenizing CI across assets leads to higher tax premium and locked-in emissions. This could indicate that on average, assets becoming locked-in solely because of the delay in taxation have lower carbon intensity than other assets. Regarding carbon budgets, reducing the reference one from Welsby et al. (2021) by 10% leaves my results unchanged. A larger 25% reduction leads to smaller overcommitted emissions in 2030. A possible explanation

is the relatively high level of tax needed in 2030 to meet the budget, largely reducing the absolute value of overcommitted emissions per asset. This reduction is not proportional to the reduction of remaining budget in 2030, with a naïve overshoot representing more than 20% of the remaining budget, much higher than in the baseline.

4.4 Carbon tax increase rate

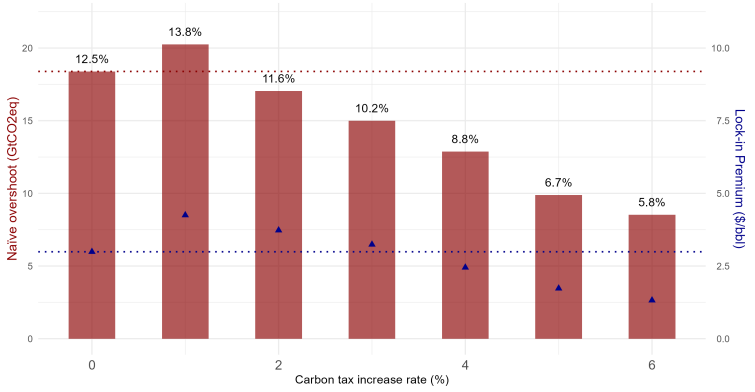

Figure 12: Sensitivity to carbon tax increase rates.

Notes: Each red bar reflects post-2030 overemissions resulting from the implementation in 2030 of the naïve carbon tax. Above each bar is indicated the corresponding share of the remaining carbon budget in 2030. The red dotted line extends the baseline results. Points correspond to the carbon tax lock-in premium, with the blue dotted line extending the baseline lock-in premium.

Although results presented in section 3 seem to indicate that the magnitude of the naïve overshoot does not differ significantly whether the policy instrument is a constant or an increasing one. However, one would expect the rate of increase to affect the results in two separate ways. On the one hand, an increasing tax implies a relatively lower tax environment in early periods, which discourage less investments as compared to the higher constant tax scenario. This translates into more similar investments realized in the baseline scenario as compared to the "No anticipation" one, reducing the cost from lock-in effects. On the other hand, the increase rate of the tax affects the quantities of emissions committed by newly developed projects: assets phased-out shortly after the delay in the constant tax scenario can continue producing longer (as long as the increasing tax is below the constant one), and other assets are phased-out earlier as the tax gets higher in later periods.

Figure 12 reports the tax lock-in premiums and naïve overshoots for carbon tax increasing at different rates, from 0% annually to 6%. Going from 0% to 1% increases the cost from lock-in effects, meaning that the second effect is positive and dominates the first

one. For all other configurations, the naïve overshoot and lock-in premium decrease as the increase rate gets higher, being more than halved for a 6% increasing tax as compared to the constant one. The first effect dominates: when the social discount rate is high, the carbon tax starts at a level low enough to only marginally reduce investments realised in new oil-producing projects.

4.5 Time-varying pass-through

The carbon tax assessed up to now was assumed to be fully borne by producers. This entails that the underlying assumption was that there is no pass-through from producers to consumers, or that this pass-through is constant over time. This assumption is relaxed in this subsection, by integrating time-varying pass-through rates into the model. I use the evolution of pass-through rates computed by Heal and Schlenker (2019) in the case of a constant tax and in the case of a 2% increase tax. While the model used in their paper is different from mine, the data on oil fields used is the same. Authors report in both cases an initial high share of the tax being supported by consumers, the cost being progressively passed on to producers over around 60 years. The initial share of the tax borne by producers is 30% in the case of a constant tax and 60% when considering a tax increasing exponentially at a 2% rate. I thus use these initial shares of the tax borne by producers, and assume they increase exponentially at a rate such that 60 years after the tax is introduced, producers bear its full cost. Equation 2 is thus adapted such that:

$$\theta_i \; Q_i^{\tau}(T) = \theta_i \; \frac{\sum_{t=T_\tau}^{T} \frac{q_{it}}{(1+\delta)^{t-T_0}} \; (1+r)^{t-2016} \; min(\Lambda_{T_\tau} \; (1+\lambda)^{t-T_\tau}, 1)}{\sum_{t=T_1}^{T} \frac{q_{it}}{(1+\delta)^{t-T_0}}}$$

with Λ_{T_τ} the initial share of the tax borne by producers when it is implemented in T_τ, and λ its increase rate such that after 60 years the tax is fully borne by producers (around 2.03% in the case of a constant tax and 1.16% for a 2% increasing one). I assume the share of tax borne by producers can not exceed 100%.

This setup differs from the case of an increasing tax for one main reason. The pass-through rate and its evolution starts *when the tax is introduced*, and not in 2016. Consequently, the level in 2030 of a delayed tax is not directly comparable to the one of a tax implemented in 2016. The share of the tax borne by producers in 2030 will be much lower in the first case as compared to the second. This adds a second dimension to the cost of

delaying taxation, since it delays the shifting of its burden to producers further into the future.

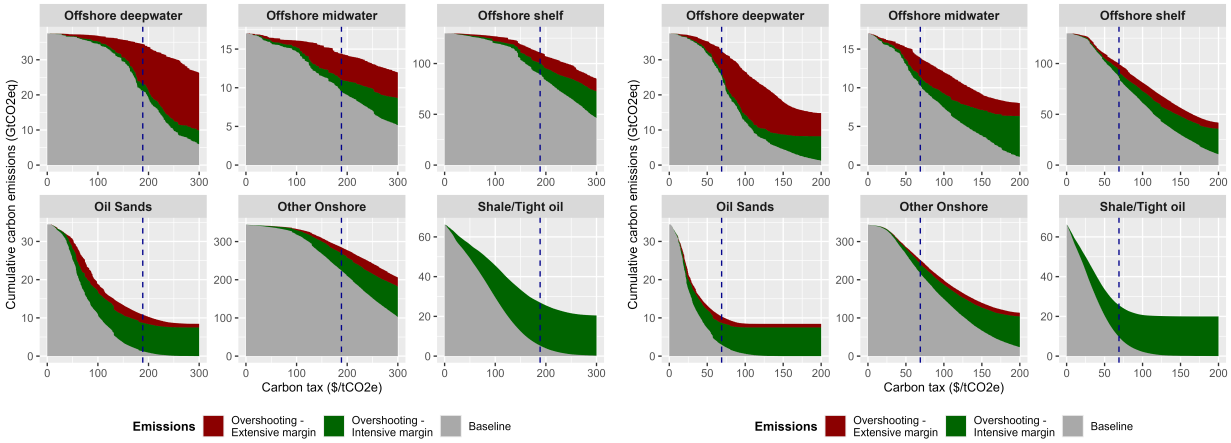

(a) Constant tax

(b) Increasing tax

Figure 13: Overshoot by supply segment accounting for time-varying pass through.

Notes: Each panel corresponds to an oil supply segment. For each panel, the x-axis is the level of carbon tax. The grey area corresponds to the "Baseline" scenario, i.e. the cumulative emissions if the tax is implemented in 2016. The green and red areas depict the overshoot resulting from delaying the tax until 2030 at the intensive and extensive margins respectively. The dotted horizontal line indicates the oil carbon budget set by Welsby et al. (2021). The dashed vertical line is the corresponding carbon tax needed in 2016 in order to constrain cumulative emissions within this budget. The scenario used to compute the overshoots is the "No anticipation" one. All computations are made assuming an oil price of $80/bbl.

Figures 13 shows the overshoot of the budget by supply segment for both the constant (panel a) and the increasing tax (panel b). Overall, overshoots are higher by around 13% at both margins compared with results presented in section 3. This is mainly explained by the fact that the passing through of the tax starts later in time, thus creating a difference in emissions committed by the same projects developed during the delay.

This translates into a higher cost stemming from the lock-in of new infrastructures. Figure 14 shows the progressive buildup of the naïve overshoot, as well as the tax premium (thus corresponding to tables 1 and 2). As expected, the levels of necessary carbon tax to meet the 1.5°C budget are particularly high. The lock-in premium is also massive, reaching 28.4$/tCO2eq in 2030 in the constant tax case. The magnitude of the premium needs to be put into perspective given the role played by the delayed passing through: whether is starts in 2016 or in 2030, the same naïve tax, equal to 265$/tCO2eq in 2030, translates into a 25.8$/tCO2eq difference in what is actually borne by producers. Accounting for this fact, the magnitude of the lock-in premium is reduced to around 3.6$/tCO2eq, consistent

with results found in other configurations. As for the naïve overshoot, its magnitude is consistently higher than in other configurations.

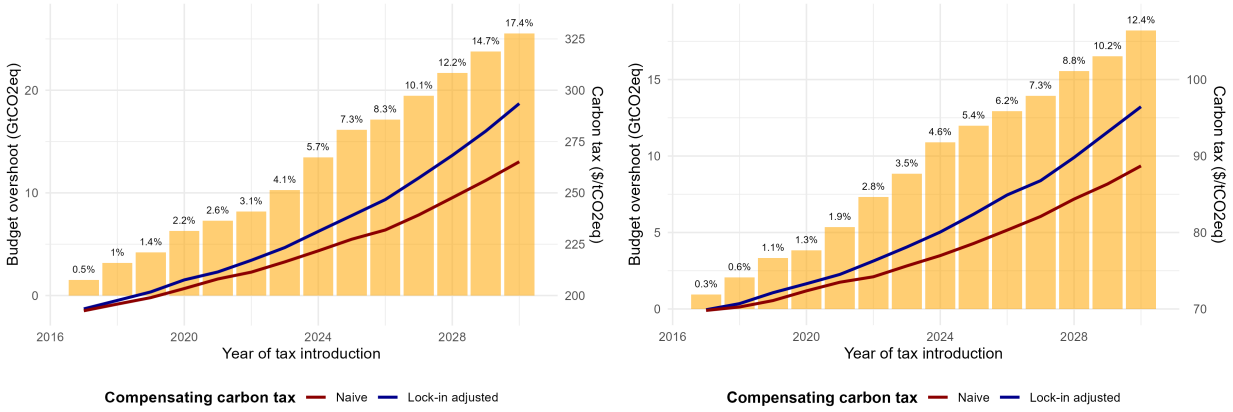

(a) Constant tax

(b) Increasing tax

Figure 14: Naive overshoot and tax premium accounting for time-varying tax pass-through.

Notes: For each year reported in x-axis, the size of bar reflects the residual overshoot resulting from implementing the naïve tax (left y-axis), with above its corresponding share of the remaining budget. The red and blue lines report the levels of the naïve and lock-in adjusted carbon taxes (right y-axis) required for each year to meet the carbon budget set by Welsby et al. (2021). The difference between the two lines corresponds to the lock-in premium. The scenario used to compute the overshoots is the "No anticipation" one. All computations are made assuming an oil price of $80/bbl.

5 Discussion and policy implications

Delaying induces more stranded assets. Findings presented in section 3 highlight the rapid pace of transition away from oil needed to meet the 1.5°C objective, and the massive cost resulting from delaying the implementation of the required ambitious climate policies. Compensating for the overproduction during the delay will translate into a vast number of stranded assets. Continued investments in new facilities further reinforces this. Among newly built assets, some will constitute stranded assets (and in turn losses for firms) when climate policies will finally be put in place. However, as shown in figure 8, a large share of these assets exhibits deep lock-in and will therefore continue producing once the tax is in place. This leads to additional stranded assets since cumulative emissions from these assets will need to be compensated through the early phase-out of already existing infrastructures.

The case for supply-side policies. Targeted policies aiming at alleviating this source of additional cost could thus be justified in terms of efficiency. The main findings of

this paper align with the literature advocating for supply-side carbon policies. As Erickson et al. (2018) argue, supply-side policies can act as complement to demand-side policies, and provide major benefits among which its direct and certain impact on fossil fuel supply and subsequent use. A major argument raised by Green and Denniss (2018) in favour of such policies is its greater potential to mobilise public support, in reason of its higher perceived benefits and distributional fairness as compared to consumption based policies. The methods used in this paper support this claim by showing the feasibility of precise accounting of expected benefits from supply-side policy, as well as the identification of potential winners and losers of the tax at the project level.

Results provide additional arguments to the growing literature supporting bans on new oil projects (Green et al., 2024), showing that targeting deepwater projects specifically could prove particularly beneficial. Indeed, worldwide carbon taxation imposed on oil producers as envisioned in the paper seems highly unlikely in the near future. At the international level, the role played by international organisations does not play in favour of supply-side policies in addressing climate change. As Van Asselt (2021) puts it, the "United Nations Framework Convention on Climate Change (UNFCCC) territorial accounting for greenhouse gas emissions emphasizes fossil fuel consumption over supply, so countries are incentivized to pursue demand-side policies to fulfil international agreements". However, this shortcoming is not the end of the road for production-based policies: a substantial share of the cost arising from infrastructure lock-in could be alleviated through targeted action.

Carbon pricing and fiscal regimes in producing countries. A further complication for supply-side policy design is the interaction between carbon pricing and existing fiscal regimes in producing countries. This paper abstracts from government take by assuming that, in the presence of a worldwide carbon tax, producing governments would renegotiate royalty structures to avoid stranding assets from which they collect a rent. Yet, producing-country governments face a Ramsey taxation problem: the optimal extraction tax balances the social return to resource rents against the production distortion it creates (Daubanes and Lasserre, 2023).

In principle, a government already at its Ramsey optimum could respond to a worldwide carbon tax by adjusting fiscal regimes so as to leave the equilibrium unchanged, substituting carbon tax revenue for royalty revenue. However, fully restoring the pre-tax equilib-

rium would in some cases require subsidising development or production. Assets that become uneconomical under the carbon tax (either because the project does not generate sufficient revenues to offset development costs or because production costs become too high towards the end of their lifetime) could only be kept active if the State offsets the carbon tax beyond what its own rent allows. Fully absorbing the tax in this way would therefore imply higher cumulative extraction and development than estimated in this paper. Conversely, contractual rigidities in existing production-sharing agreements and concession contracts may prevent renegotiation in the short term, causing the carbon tax to pile on top of the existing fiscal regime and further constrain investment and production beyond what the model predicts.

The assumption made in this paper implicitly lies between these two extremes: governments are assumed to absorb the carbon tax to the extent necessary to keep viable assets in production (potentially capturing a higher rent in the process) without subsidising uneconomical ones, so that assets not resilient to the carbon tax are either phased out earlier than in a no-tax world or simply remain undeveloped. Under this assumption, the overcommitted emissions identified in the paper represent an intermediate scenario, with the full-absorption case providing an upper bound and the full-rigidity case a lower bound on the true extent of infrastructure lock-in. These interactions further suggest that supply-side bans on new field development may in some cases be a more administratively tractable instrument than relying on carbon pricing to overcome the fiscal inertia of producing States.

Targeted policies. My results show that carbon lock-in is much deeper in assets concentrated in specific regions and/or using specific extracting technologies, making potential national or sub-national actions effective at mitigating the cost of delaying global action. Figure 15 shows post-2030 overcommitted emissions, aggregated by province, resilient to the 2030 2% increasing naïve tax, originating from investments that would not have been carried out should the 57.1$/tCO2eq tax be implemented in 2016. With this restrictive definition, the total 15 GtCO2eq overcommitted emissions originate from 548 assets scattered around the globe. Despite large over-investments in the past 10 years, future investments over the next five years are crucial. As highlighted in the figure, the cost from carbon lock-in has yet to become fully locked-in in many regions. In terms of distribution, carbon lock-in effects occur in most oil producing countries, with offshore projects representing more than 75% of the total (half of which comes from deepwater projects). Identifying specific assets re-

sponsible for it, I find that it is highly concentrated, with only 31 assets making up for 50% of the total. Deepwater offshore projects play a major role: assets along the Brazilian and Guyanese coasts, as well as in the Gulf of Mexico, represent alone more than 30% of total overcommitted emissions. Issuing targeted bans towards these problematic projects could go a long way in avoiding some of the cost arising from infrastructure lock-in.

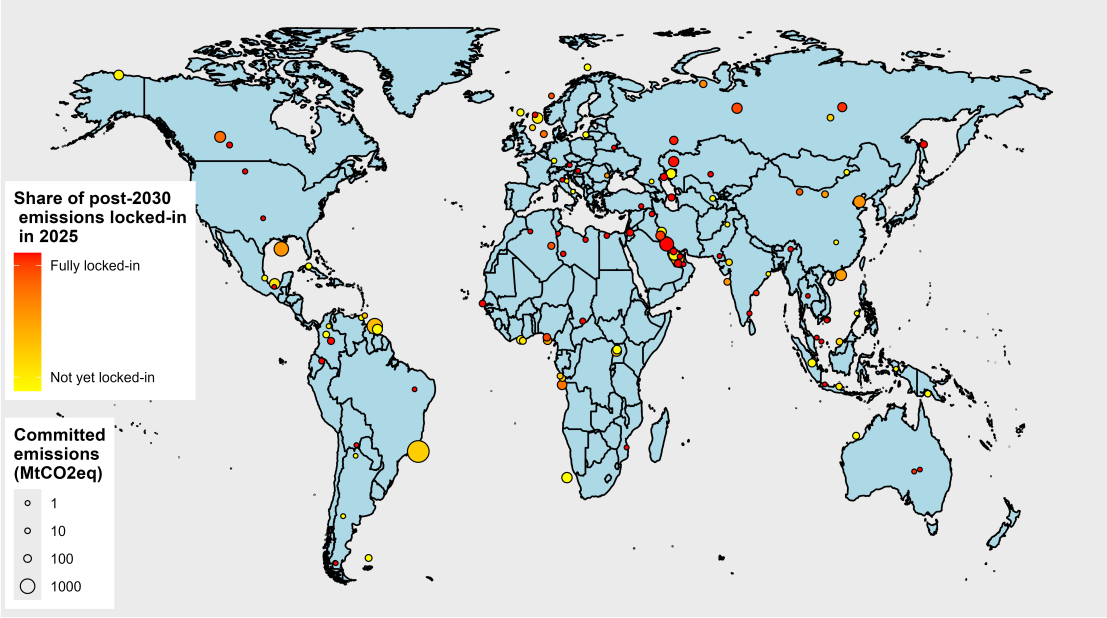

Figure 15: Overcommitted emissions aggregated by province (increasing tax).

Notes: Each circle correspond to one province in which overcommitted emissions exist in the "no anticipation" scenario. A province is a level of disaggregation available in Rystad's data situated between the asset-level and the country-level. Circle sizes depend on the quantities overcommitted and colors depend on the degree of commitment (share of the reserves committed) of the associated quantities as of 2025. Committed emissions from a given asset are attributed to the year from which the full development of this asset becomes inevitable. The level in 2016 of the delayed tax used to identify projects responsible for overcommitted emissions is 57.1$/tCO2eq. The level of 2030 tax used to compute the quantities of overcommitted emissions per project is 66.3$/tCO2eq. The scenario used for this graph in the "no anticipation" one. All computations are made assuming a reference Brent oil price of $80/bbl.

The effectiveness of public action. Convincing large producing countries to issue such bans might prove difficult, with OPEC countries being the leading opposition to any moratorium on fossil fuel extraction and new field development (Lujala et al., 2022), but smaller initiatives could prove effective. As Piggot (2018) and Carter and McKenzie (2020) explain, most bans issued thus far have originated from public mobilisation efforts focused on specific projects. These targeted initiatives have demonstrated higher ability of public action to more directly influence political decisions. Given the important concentration of overcommitted emissions in a small number of projects, actions by the public directed to specific projects could have tremendous impact. Among the more recent successful public

mobilisations is the definitive cancelling of the Keystone XXL pipeline. Erickson and Lazarus (2014) previously estimated that the building of the pipeline could embed up to 110 MtCO2eq of annual emissions in reason of increased Canadian oil sands production. On the opposite end, an unsuccessful instance of public initiative can be found with the Alaskan Willow project, extensively discussed in 2022 in the American public debate. Finally approved by the Biden administration in 2023, this project embeds 238 MtCO2eq of overcommitted emissions over the next 25 years, resilient to the higher levels of tax required in 2030 to meet the carbon budget. It is responsible for 1.4% of the 2030 naïve overshoot computed in table 2.

Public action on larger oil projects could prove even more effective: the deepwater offshore Buzios project in Brazil concentrates 8.3% of the naïve overshoot. More than 60% of deepwater projects developed during the delay are majoritarily owned and operated by majors (BP, Chevron, Eni, ExxonMobil, Shell and TotalEnergies), and their development is often heavily dependent on technologies mastered by these firms. Developed countries could impose a ban on these specific projects, and incentivise major firms to redirect the important financial means required to develop these projects towards cleaner development opportunities for the corresponding countries. This leverage is further reinforced by the results from section 3: screening projects with the full carbon tax reduces the overshoot at the extensive margin by a striking 84% for the subset of assets with publicly-listed firm shareholders, highlighting the potential of investor pressure as a complementary channel.

Energy transition and equity. Results also provide insights on the path towards an equitable and fair transition. As explained by Pye et al. (2020), "prioritising equity criteria could result in high-cost reserves and those that require new infrastructure being developed, while low-cost reserves and those with infrastructure in place are phased out". Focusing on the key sources of overshoot at the intensive and extensive margins could allow to resolve this tension between equity and cost-efficiency.

On the one hand, overproduction during the delay is mainly due to North American production, in particular from shale and oil sands. Early action from the more developed countries could be justified in terms of global distributive justice (Lenferna, 2018), and would reap important returns in terms of carbon emissions abatement. Localised actions in specific areas (e.g. Alberta shale and oil sands in Canada, Texas and North Dakota shale in

the US) would lead to a significant reduction in overemissions during the delay while being constrained to limited jurisdictions. On the other hand, the deep lock-in concentrated in offshore projects largely owned and operated by major firms headquartered in developed countries creates a complementary lever: redirecting the considerable financial resources currently committed to these projects towards renewable development in oil-dependent countries would address both the efficiency and the distributional dimensions of the transition simultaneously.

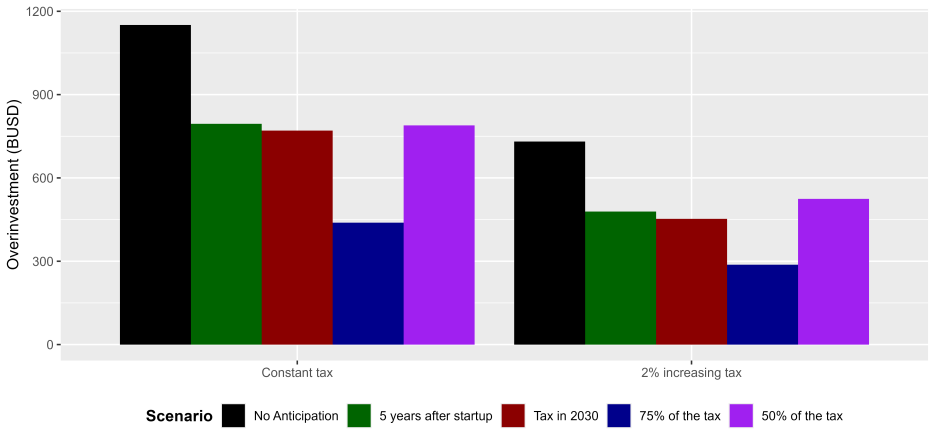

Figure 16: 2025 discounted total over-investment by scenario.

Notes: Each bar reports the 2025 discounted sum (2% discount rate) of investments in oil projects dependent on the tax being delayed. The level of 2016 delayed tax used to identify over-investments is 99.3$/tCO2eq. All computations are made assuming a reference Brent oil price of $80/bbl.

The opportunity cost of lock-in. Not only does over-investment caused by the delay result in more stranded assets, and in turn losses for oil and gas firms as well as governments, but it also uses considerable financial resources which could be otherwise used to develop renewable energy. Figure 16 reports the sum of these inefficient investments carried out because of the delay, discounted as of 2025. In the case of a constant tax, over-investments represent more than 80 BUSD annually, for a total of 1150 BUSD. For reference, this level of annual inefficient investments worldwide is equivalent to 75% of total investments in renewables in Europe in 2023 (International Energy Agency, 2024). In the case of an increasing tax, the carbon tax is lower during the delay, making more assets economically viable even in presence of the tax. Total over-investment is reduced but remains considerable, amounting to 730 BUSD. Alternative sets of anticipations by investors can reduce this figure by up to 60% (scenario of tax underestimated by 25%). Nonetheless, directing these financial

resources towards building renewables capacities, especially in oil-dependent countries, could greatly accelerate the pace of the transition.

6 Conclusion

This study quantifies the climate consequences of delaying carbon taxation in the oil sector, revealing that postponing the implementation of a constant 99.3$/tCO2eq tax (respectively a 2% increasing 57.1$/tCO2eq), aligned with the 1.5°C target, from 2016 to 2030 results in a 33% overshoot (respectively 19%) of the oil carbon budget. The overshoot arises from two distinct mechanisms: on the one hand, overproduction occurs during the delay, concentrated in high-cost, high-emission assets such as Canadian oil sands and U.S. shale. On the other hand, long-term infrastructure lock-in builds up, driven by capital-intensive offshore projects that are costly to develop but cheap to operate once developed. Over-committed emissions from locked-in assets are highly concentrated, with deepwater projects along the coasts of Brazil, Guyana and in the Gulf of Mexico contributing disproportionately. These findings entail important policy implications. While global supply-side climate policies face significant opposition, my results indicate that targeted interventions focusing on specific technologies (such as deepwater projects) or specific regions could yield substantial emissions reductions. Strategies aiming at counteracting the key sources of over-emissions can be aligned with fairness and equity considerations in transitioning away from fossil fuels.

The analysis introduces the concept of carbon tax lock-in premium: the additional tax premium required to offset the downward shifts in the supply curve due to cumulative investments in new production facilities. Ignoring these lock-in effects leads policymakers to systematically underprice carbon, despite correctly defining the remaining carbon budget. For a constant tax as well as an increasing one, implementing a naïve tax in 2030 would result in a 17 to 19 GtCO2eq overshoot of the budget, originating in emissions taking place after 2030. Crucially, this residual overshoot tends to buildup at a rapid pace: failing to account for only the next five years of cumulative investments in new production facilities would result in overshooting the carbon budget by over 10 GtCO2eq. Even short-term delay in policy action translates into costly long-term commitments. By highlighting the irreversible nature of infrastructure lock-in and the cumulative cost of policy delays, this study highlights

the necessity of supply-side climate policies, not designed as a substitute for carbon pricing, but as a complementary tool to mitigate long-term carbon lock-in effects.

Several limitations of this analysis point to directions for future research. First, the paper focuses exclusively on the oil sector, and the results should be interpreted under the implicit assumption that comparable policies are applied simultaneously to gas and coal. To the extent that a carbon tax on oil alone induces substitution toward other fossil fuels, the budget overshoots documented here represent a lower bound on the true cost of delay. Second, the model abstracts from the Green Paradox mechanism: producers anticipating future taxation may be incentivised to invest earlier in new producing facilities, thereby amplifying overproduction during the delay and reinforcing lock-in. The results should therefore be read as conservative estimates of the full cost of policy inaction. Additionally, the model abstracts from the Ramsey taxation problem faced by resource-owning governments: in practice, producing States optimise their fiscal regime to capture extraction rents, and the interaction between this optimal fiscal structure and an externally imposed carbon price is non-trivial. Exploring the joint determination of optimal resource taxation and carbon pricing in a world of delayed policy remains an important direction for future work. Finally, while the model quantifies the long-term cost of delayed policy implementation, it does not endogenize the political economy mechanisms that generate such delays. Findings suggest that ignoring lock-in effects likely leads to underestimating the costs of postponing the carbon policy. Future research could usefully compare the magnitude of the additional cost induced by carbon lock-in to the political constraints explaining the limited ambition of currently observed climate policies.

References

Abella, J. P. and Bergerson, J. A. (2012). Model to investigate energy and greenhouse gas emissions implications of refining petroleum: Impacts of crude quality and refinery configuration. *Environmental science & technology*, 46(24):13037–13047.

Ahluwalia, M. B. (2017). The business of pricing carbon: how companies are pricing carbon to mitigate risks and prepare for a low-carbon future. *Center for Climate and Energy Solutions, Arlington, VA*.

Ahlvik, L., Andersen, J. J., Hamang, J. H., and Harding, T. (2022). Quantifying supply-side climate policies.

Ahlvik, L. and Harding, T. (2025). Resource investments and the timing of tax deductions. *Journal of the Association of Environmental and Resource Economists*, 12(1):1–32.

Anderson, S. T., Kellogg, R., and Salant, S. W. (2018). Hotelling under pressure. *Journal of Political Economy*, 126(3):984–1026.

Bailey, W., Couët, B., Lamb, F., Simpson, G., and Rose, P. (2000). Taking a calculated risk. *Oilfield Review*, 12(3):20–35.

Berntsen, M., Bøe, K. S., Jordal, T., and Molnár, P. (2018). Determinants of oil and gas investments on the norwegian continental shelf. *Energy*, 148:904–914.

Bertram, C., Johnson, N., Luderer, G., Riahi, K., Isaac, M., and Eom, J. (2015). Carbon lock-in through capital stock inertia associated with weak near-term climate policies. *Technological Forecasting and Social Change*, 90:62–72.

Cairns, R. D. and Smith, J. L. (2019). The Green Paradox, A Hotelling Cul de Sac. *Economics of Energy & Environmental Policy*, 8(2).

Carter, A. V. and McKenzie, J. (2020). Amplifying "Keep It in the Ground" First-Movers: Toward a Comparative Framework. *Society & Natural Resources*, 33(11):1339–1358.

Climate Watch (2023). Ghg emissions. *World Resources Institute: Washington, DC, USA*.

Daubanes, J. X. and Lasserre, P. (2023). How should the use of nonrenewables be taxed under a public budget constraint? *Resource and Energy Economics*, 73:101375.

Davis, S. J. and Socolow, R. H. (2014). Commitment accounting of CO_2 emissions. *Environmental Research Letters*, 9(8):084018.

Dietz, S. and Venmans, F. (2019). Cumulative carbon emissions and economic policy: in search of general principles. *Journal of Environmental Economics and Management*, 96:108–129.

Elvidge, C. D., Zhizhin, M., Baugh, K., Hsu, F.-C., and Ghosh, T. (2015). Methods for global survey of natural gas flaring from visible infrared imaging radiometer suite data. *Energies*, 9(1):14.

Eom, J., Edmonds, J., Krey, V., Johnson, N., Longden, T., Luderer, G., Riahi, K., and Van Vuuren, D. P. (2015). The impact of near-term climate policy choices on technology and emission transition pathways. *Technological Forecasting and Social Change*, 90:73–88.

Erickson, P., Down, A., Lazarus, M., and Koplow, D. (2017). Effect of subsidies to fossil fuel companies on United States crude oil production. *Nature Energy*, 2(11):891.

Erickson, P. and Lazarus, M. (2014). Impact of the keystone xl pipeline on global oil markets and greenhouse gas emissions. *Nature Climate Change*, 4(9):778–781.

Erickson, P. and Lazarus, M. (2018). Would constraining US fossil fuel production affect global CO2 emissions? A case study of US leasing policy. *Climatic Change*, 150(1-2):29–42.

Erickson, P., Lazarus, M., and Piggot, G. (2018). Limiting fossil fuel production as the next big step in climate policy. *Nature Climate Change*, 8(12):1037–1043.

Fæhn, T., Hagem, C., Lindholt, L., Mæland, S., and Rosendahl, K. E. (2017). Climate policies in a fossil fuel producing country–demand versus supply side policies. *The Energy Journal*, 38(1).

Fattouh, B. (2007). *OPEC pricing power: The need for a new perspective*. Oxford Institute for Energy Studies.

Fisch-Romito, V., Guivarch, C., Creutzig, F., Minx, J. C., and Callaghan, M. W. (2021). Systematic map of the literature on carbon lock-in induced by long-lived capital. *Environmental Research Letters*, 16(5):053004.

Fried, S., Novan, K., and Peterman, W. B. (2021). The Macro Effects of Climate Policy Uncertainty. *Finance and Economics Discussion Series*, 2021(015):1–50.

Ghosh, S., Jean, L., and Saha, P. (2025). The last good owner effect: Ownership structure and environmental outcomes in oil and gas.

Gollier, C. (2009). *Pricing the Planet's Future: The Economics of Discounting in an Uncertain World*. Princeton University Press.

Green, F., Bois von Kursk, O., Muttitt, G., and Pye, S. (2024). No new fossil fuel projects: The norm we need. *Science*, 384(6699):954–957.

Green, F. and Denniss, R. (2018). Cutting with both arms of the scissors: the economic and political case for restrictive supply-side climate policies. *Climatic Change*, 150(1-2):73–87.

Greenstone, M., Kopits, E., and Wolverton, A. (2013). Developing a social cost of carbon for us regulatory analysis: A methodology and interpretation. *Review of environmental economics and policy*.

Grubb, M., Wieners, C., and Yang, P. (2021). Modeling myths: On DICE and dynamic realism in integrated assessment models of climate change mitigation. *WIREs Climate Change*, 12(3).

Heal, G. and Schlenker, W. (2019). Coase, Hotelling and Pigou: The Incidence of a Carbon Tax and CO Emissions. Technical Report w26086, National Bureau of Economic Research, Cambridge, MA.

Höök, M. and Tang, X. (2013). Depletion of fossil fuels and anthropogenic climate change—a review. *Energy policy*, 52:797–809.

Insley, M. (2017). Resource extraction with a carbon tax and regime switching prices: Exercising your options. *Energy Economics*, 67:1–16.

International Energy Agency (2024). World energy investment 2024. Accessed 2026-02-19; includes regional and European Union renewable energy investment data.

International Monetary Fund (2024). Regional Economic Outlook: Middle East and Central Asia, April 2024: An Uneven Recovery amid High Uncertainty. International Monetary Fund. Working assumptions include Brent crude oil price averaging $78.61/bbl in 2024 and $73.68/bbl in 2025.

IPCC (2006). *2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 2: Energy*. Institute for Global Environmental Strategies (IGES), Hayama, Japan.

Jahn, F., Cook, M., and Graham, M. (2008). The field life cycle. *Developments in Petroleum Science*, 55:1–7.

Jakob, M., Luderer, G., Steckel, J., Tavoni, M., and Monjon, S. (2012). Time to act now? Assessing the costs of delaying climate measures and benefits of early action. *Climatic Change*, 114(1):79–99.

Kellogg, R. (2024). The end of oil. Technical report, National Bureau of Economic Research.

Kotlikoff, L., Kubler, F., Polbin, A., Sachs, J., and Scheidegger, S. (2021). MAKING CARBON TAXATION A GENERATIONAL WIN WIN. *International Economic Review*, 62(1):3–46.

Lazarus, M., Erickson, P., and Tempest, K. (2015). Carbon lock-in from fossil fuel supply infrastructure.

Lazarus, M. and van Asselt, H. (2018). Fossil fuel supply and climate policy: exploring the road less taken. *Climatic Change*, 150(1-2):1–13.

Lemoine, D. and Rudik, I. (2017). Steering the climate system: using inertia to lower the cost of policy. *American Economic Review*, 107(10):2947–2957.

Lenferna, G. A. (2018). Can we equitably manage the end of the fossil fuel era? *Energy Research & Social Science*, 35:217–223.

Luderer, G., Pietzcker, R. C., Bertram, C., Kriegler, E., Meinshausen, M., and Edenhofer, O. (2013). Economic mitigation challenges: how further delay closes the door for achieving climate targets. *Environmental Research Letters*, 8(3):034033.

Lujala, P., Le Billon, P., and Gaulin, N. (2022). Phasing out fossil fuels: Determinants of production cuts and implications for an international agreement. *Global Environmental Politics*, 22(4):95–128.

Masnadi, M. S., El-Houjeiri, H. M., Schunack, D., Li, Y., Englander, J. G., Badahdah, A., Monfort, J.-C., Anderson, J. E., Wallington, T. J., Bergerson, J. A., et al. (2018). Global carbon intensity of crude oil production. *Science*, 361(6405):851–853.

McGlade, C. and Ekins, P. (2015). The geographical distribution of fossil fuels unused when limiting global warming to 2 °C. *Nature*, 517(7533):187–190.

Metcalf, G. E. (2019). On the economics of a carbon tax for the united states. *Brookings Papers on Economic Activity*, 2019(1):405–484.

Nordhaus, W. D. (2007). A review of the stern review on the economics of climate change. *Journal of Economic Literature*, 45(3):686–702.

Pierru, A., Smith, J., and Zamrik, T. (2018). Opec's impact on oil price volatility: The role of spare capacity. *The Energy Journal*, 39(2):173–196.

Piggot, G. (2018). The influence of social movements on policies that constrain fossil fuel supply. *Climate Policy*, 18(7):942–954.

Pye, S., Bradley, S., Hughes, N., Price, J., Welsby, D., and Ekins, P. (2020). An equitable redistribution of unburnable carbon. *Nature Communications*, 11(1):3968.

Riahi, K., Kriegler, E., Johnson, N., Bertram, C., den Elzen, M., Eom, J., Schaeffer, M., Edmonds, J., Isaac, M., Krey, V., Longden, T., Luderer, G., Méjean, A., McCollum, D. L., Mima, S., Turton, H., van Vuuren, D. P., Wada, K., Bosetti, V., Capros, P., Criqui, P., Hamdi-Cherif, M., Kainuma, M., and Edenhofer, O. (2015). Locked into Copenhagen pledges — Implications of short-term emission targets for the cost and feasibility of long-term climate goals. *Technological Forecasting and Social Change*, 90:8–23.

Seto, K. C., Davis, S. J., Mitchell, R. B., Stokes, E. C., Unruh, G., and Ürge Vorsatz, D. (2016). Carbon Lock-In: Types, Causes, and Policy Implications. *Annual Review of Environment and Resources*, 41(1):425–452.

Sinn, H.-W. (2008). Public policies against global warming: a supply side approach. *International Tax and Public Finance*, 15(4):360–394.

Smith, C. J., Forster, P. M., Allen, M., Fuglestvedt, J., Millar, R. J., Rogelj, J., and Zickfeld, K. (2019). Current fossil fuel infrastructure does not yet commit us to 1.5 °C warming. *Nature Communications*, 10(1):101.

Stefanski, R., Ahlvik, L., Andersen, J. J., Harding, T., and Trew, A. (2025). Extracting wedges: Misallocation and taxation in the oil industry. Technical report, CESifo Working Paper.

Stern, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge University Press.

Tang, B.-J., Zhou, H.-L., Chen, H., Wang, K., and Cao, H. (2017). Investment opportunity in china's overseas oil project: An empirical analysis based on real option approach. *Energy Policy*, 105:17–26.

Tong, D., Zhang, Q., Zheng, Y., Caldeira, K., Shearer, C., Hong, C., Qin, Y., and Davis, S. J. (2019). Committed emissions from existing energy infrastructure jeopardize 1.5 °C climate target. *Nature*, 572(7769):373–377.

Trout, K., Muttitt, G., Lafleur, D., Van De Graaf, T., Mendelevitch, R., Mei, L., and Meinshausen, M. (2022). Existing fossil fuel extraction would warm the world beyond 1.5 °C. *Environmental Research Letters*, 17(6):064010.

Unruh, G. C. (2000). Understanding carbon lock-in. *Energy Policy*, 28(12):817–830.

Van Asselt, H. (2021). Governing fossil fuel production in the age of climate disruption: Towards an international law of 'leaving it in the ground'. *Earth System Governance*, 9:100118.

van der Ploeg, F. and Withagen, C. (2015). Global Warming and the Green Paradox: A Review of Adverse Effects of Climate Policies. *Review of Environmental Economics and Policy*, 9(2):285–303.

Van Vuuren, D. P., Van Soest, H., Riahi, K., Clarke, L., Krey, V., Kriegler, E., Rogelj, J., Schaeffer, M., and Tavoni, M. (2016). Carbon budgets and energy transition pathways. *Environmental Research Letters*, 11(7):075002.

Venables, A. J. (2014). Depletion and Development: Natural Resource Supply with Endogenous Field Opening. *Journal of the Association of Environmental and Resource Economists*, 1(3):313–336.

Wachtmeister, H. and Höök, M. (2020). Investment and production dynamics of conventional oil and unconventional tight oil: Implications for oil markets and climate strategies. *Energy and Climate Change*, 1:100010.

Weitzman, M. L. (1998). Why the far-distant future should be discounted at its lowest possible rate. *Journal of Environmental Economics and Management*, 36:201–208.

Weitzman, M. L. (2007). A review of the stern review on the economics of climate change. *Journal of Economic Literature*, 45(3):703–724.

Welsby, D., Price, J., Pye, S., and Ekins, P. (2021). Unextractable fossil fuels in a 1.5 °C world. *Nature*, 597(7875):230–234.

Wood Mackenzie (2013). Impacts of Delaying IDC Deductibility (2014-2025).

A Appendix figures

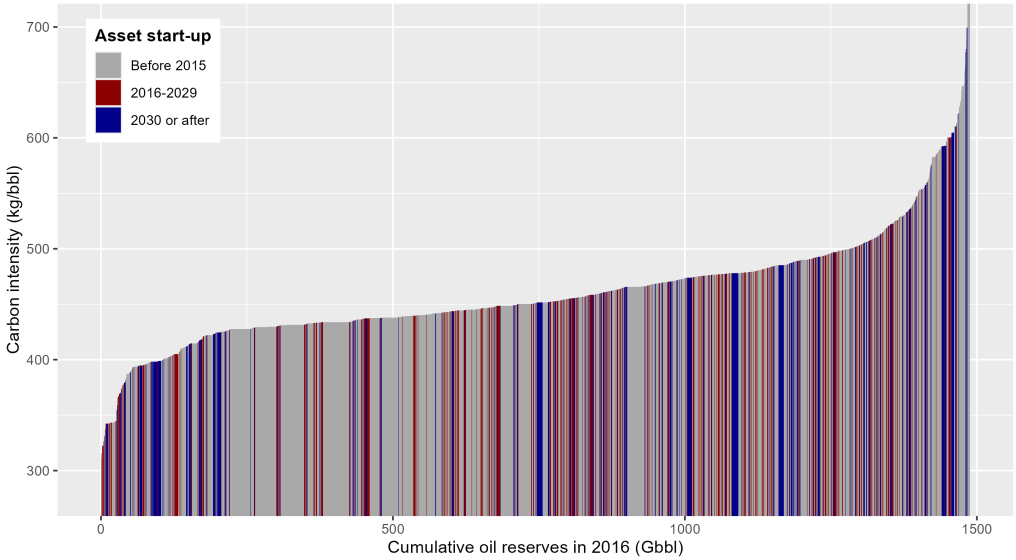

Figure 17: Cumulative oil reserves in 2016 ordered by carbon intensities.

Notes: The y-axis corresponds to lifecycle carbon intensity, thus including upstream, midstream and downstream emissions. Cumulative reserves are split between assets whose production started before 2016, assets developed during the delay (2016-2029), and assets yet-to-be developed in 2030.

Economics Category	The Economics Category splits the Economics variable (=Revenue) into higher granularity than Economics Group, e.g. capital expenses is subdivided into facility and well expenses. The sum of all items equals the revenue. Use Economics Category to analyze breakdown of cost models or specific cost elements.
Free Cash Flow	FCF is the incurred cash flow available to investors and creditors, thus largely used in DCF valuation analysis. FCF = revenue-capex-opex-gvt take. FCF can be compared with cash from operations and investments activity in the cash flow statement from companies (only upstream). The present value of (all discounted) FCF equals NPV (Net Present Value also called Gross Asset Value) that can be compared with Enterprise Value (=Market cap+Debt for upstream activities), or compared with asset transaction costs. The discount factor in UCube is 10% flat for all assets.
Government Profit Oil	Gvt PO is the PSA equivalent to petroleum taxes, but paid in kind. Gvt PO reduces the company's entitelement production and is thus treated as a royalty effect in company reports.
Royalty effects	Royalty effects is the sum of all gross taxes, including royalty and oil and export duties.
Income Tax	Income tax is a sum of all profit based taxes, ie corporate taxes, where the tax rate is equal to the country's corporate tax rate, and special petroleum taxes. The tax base is the net sales revenue minus opex and depreciation. For petroleum taxes there may be applied an additional deduction called uplift. In many cases, some profit based taxes can be deducted in the tax base for other profit taxes. The income taxes in the UCube assumes there is no interest expenses.
Bonuses	Signature bonus for the license at award, as officially reported.
Well Capex	Well capex is capitalized costs related to well construction, including drilling costs, rig lease, well completion, well stimulation, steel costs and materials.
Facility Capex	Development capex except well construction costs, includes costs to develop, install, maintain and modify surface installations and infrastructure. As reported by operators, field partners or officials, or modelled. Distribution of costs over time is largely modelled, one exception is Norway, where all historic costs are as reported in national budgets.
Exploration Capex	*Costs incurred to find and prove hydrocarbons: seismic, wildcat and appraisal wells, general engineering costs, based on reports and budgets or modelled.*
Abandonment cost	Abandonment cost represents the costs associated with shutting down and dismantling the surface and subsea facilities such as surface equipment, platforms, subsea structures and pipelines. Activities included is topside and jacket removal as well as pipeline cleaning and decommissioning, onshore disposal, surveys and monitoring, and project management costs. Note that the Abandonment cost only includes the remaining P&A work for wells, and not full cycle P&A work. Abandonment cost is mostly modelled and assumed to be incurred and expensed over couple of years after the production has ceased. The estimate is proportional to initial greenfield capex.
Production Opex	Represents operational expenses directly related to the production activity. The category includes materials, tools, maintenance, equipment lease costs and operation related salaries. Depreciation and other non-cash items are not included.
Transportation Opex	Represents the costs of bring the oil and gas from the production site/processing plant to the pricing point (only upstream transportation). The category includes transport fees and blending costs.
SG&A Opex	Represent operation expenses not directly associated with field operations. The category includes administrative staff costs, office leases, related benefits (stocks and stock option plans) and professional expenses (legal, consulting, insurance). Only E&P related SG&A are included.
Taxes Opex	Represent gross tax reported by companies under operational costs and applies for US and Brazil. For US the category represent ad valorem taxes (county based) and severance taxes (state based).

Figure 18: **Rystad economic data - Split items**: Descriptions of the data are provided by Rystad alongside with the UCube database.

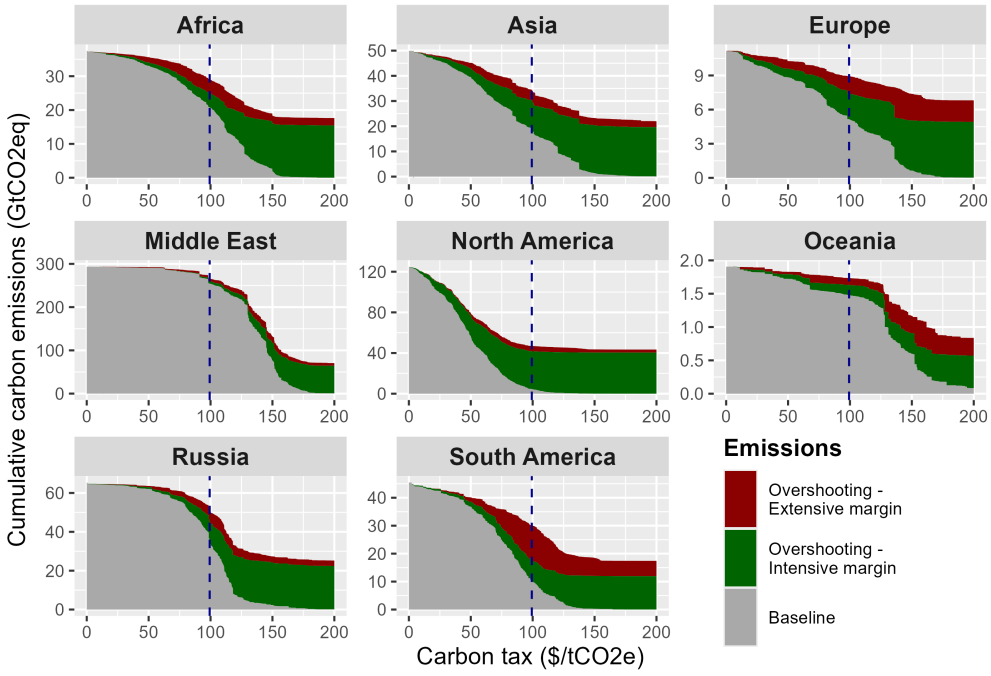

Figure 19: Oil carbon budget overshooting by world region.

Notes: Each panel corresponds to a region of the world. For each panel, the x-axis is the level of carbon tax. The grey area corresponds to the baseline scenario, i.e. the cumulative emissions if the tax is implemented in 2016. The green and red areas depicts the overshoot resulting from delaying the tax until 2030 at the intensive and extensive margins respectively. The dotted horizontal line indicates the oil carbon budget set by Welsby et al. (2021). The dashed vertical line is the corresponding carbon tax needed in 2016 in order to constrain cumulative emissions within this budget. The scenario used to compute the overshoots is the "no anticipation" one. All computations are made assuming an oil price of $80/bbl.

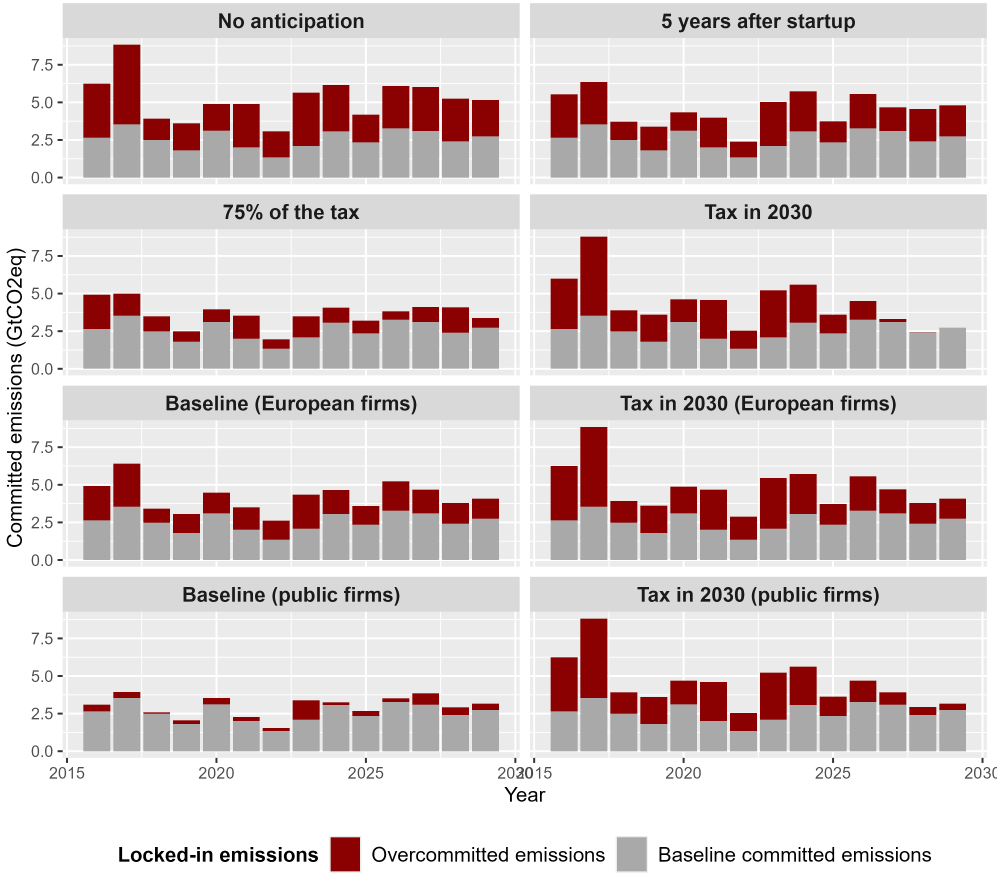

Figure 20: Yearly committed emissions by scenario.

Notes: Each panel corresponds to a scenario. Each bar corresponds to the amount of committed emissions from investments in the development of new oil assets. The grey share of the bar depicts committed emissions that are emitted in the baseline scenario. The red share corresponds to committed emissions from assets that are not developed in the baseline scenario, thus representing the cost arising from carbon lock-in effects. Committed emissions from a given asset are attributed to the year from which the full development of this asset becomes inevitable, which depends on the level of tax. The level of tax used for this graph is \$100/tCO2eq. All computations are made assuming a reference Brent oil price of \$80/bbl.

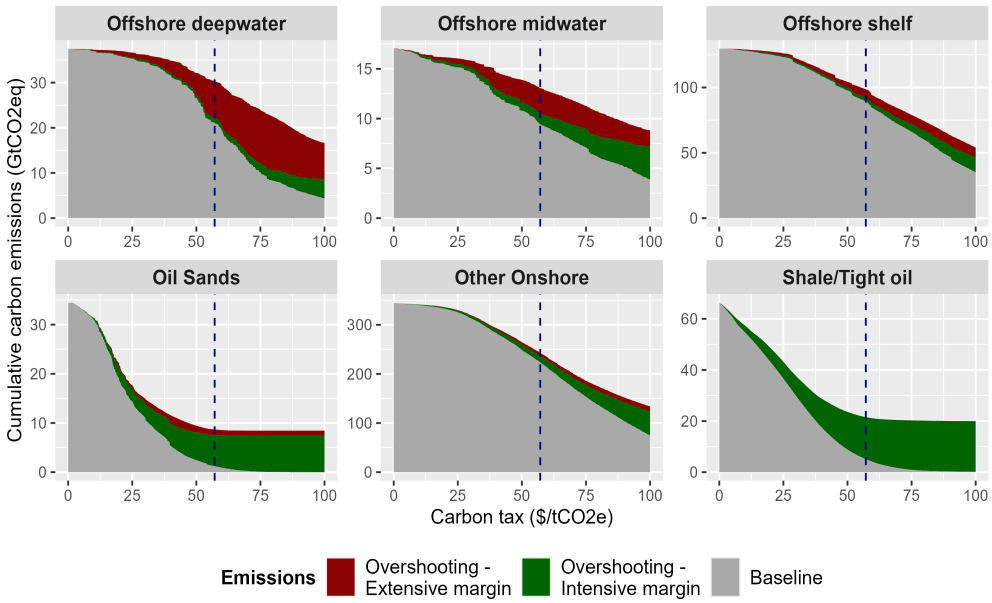

Figure 21: Oil carbon budget overshooting by supply segment.

Notes: Each panel corresponds to an oil supply segment. For each panel, the x-axis is the level of carbon tax. The grey area corresponds to the baseline scenario, i.e. the cumulative emissions if the tax is implemented in 2016. The green and red areas depicts the overshoot resulting from delaying the tax until 2030 at the intensive and extensive margins respectively. The dotted horizontal line indicates the oil carbon budget set by Welsby et al. (2021). The dashed vertical line is the corresponding carbon tax needed in 2016 in order to constrain cumulative emissions within this budget. The scenario used to compute the overshoots is the "no anticipation" one. All computations are made assuming an oil price of $80/bbl.

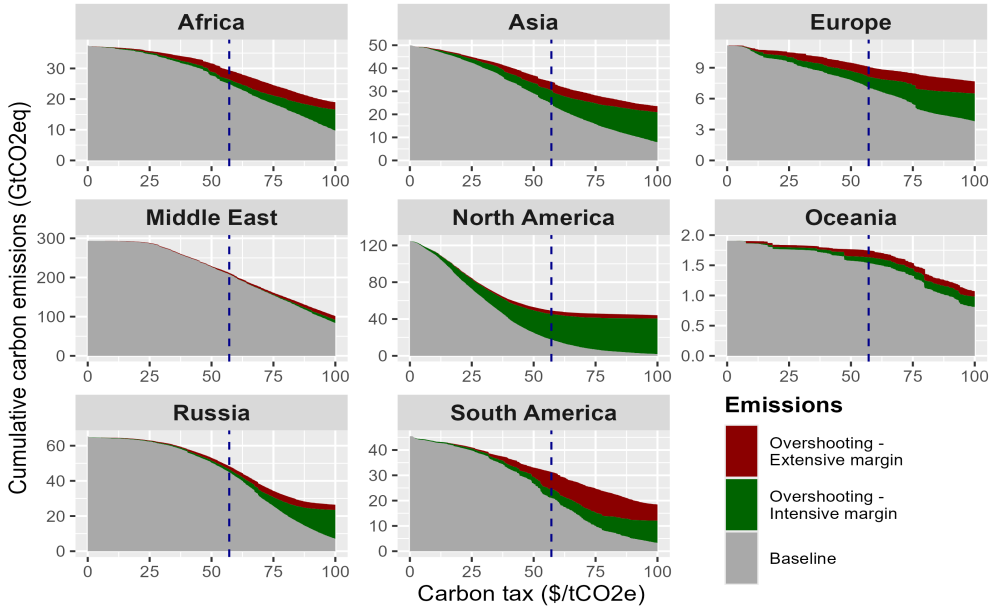

Figure 22: Oil carbon budget overshooting by supply segment.

Notes: Each panel corresponds to an oil supply segment. For each panel, the x-axis is the level of carbon tax. The grey area corresponds to the baseline scenario, i.e. the cumulative emissions if the tax is implemented in 2016. The green and red areas depicts the overshoot resulting from delaying the tax until 2030 at the intensive and extensive margins respectively. The dotted horizontal line indicates the oil carbon budget set by Welsby et al. (2021). The dashed vertical line is the corresponding carbon tax needed in 2016 in order to constrain cumulative emissions within this budget. The scenario used to compute the overshoots is the "no anticipation" one. All computations are made assuming an oil price of $80/bbl.

Oil Type	Average discount	2016 reserves (Gbbl)
Regular	-4.5%	919
NGL	-50.6%	253
Light	-2.1%	248
Condensate	-5.7%	120
Sour (> 3%)	-4.4%	55
Bitumen	-29%	47
Heavy Oil 15–19	-14%	34
Heavy Oil 20–22	-10.4%	22
Extra Heavy Oil	-12.3%	21
Synthetic crude	-7.4%	19
Total	-12%	1738
excluding NGL	-5.4%	1485

Table 3: Discounts relative to Brent oil prices by oil type.

Notes: The table reports the reserve weighted average of discounts relative to the reference Brent oil price for each specific oil type. Discounts are computed as average discounts observed and predicted by Rystad Energy over the 2016-2100 period.

B Field level carbon intensities

While Rystad's databases provide detailed asset-level data on oil production, they lack comprehensive carbon intensity (CI) estimates. This appendix fills this gap by combining engineering-based models (OPGEE, PRELIM, OPEM), satellite-derived flaring data (VIIRS), and predictive regressions to estimate lifecycle CI for all assets, including those not yet in production. Carbon emissions related to oil can be split into 3 stages (Gordon et al., 2015). The upstream stage includes all emissions arising from extraction processes to the transport of crude oil to the refinery inlet. The midstream stage corresponds to refining processes and are heavily dependent on the refinery configuration. Finally, the downstream stage comprises emissions from the transport of oil products to their end-user and its final use. While downstream emissions usually represents around 80% of total emissions, the upstream stage features important variability from one field to the other.

For each of these stage, I rely on a state-of-the-art model developed by the Oil-Climate Index on a sample of currently produced oil to estimate their carbon intensity. I then fit a regression model for each stage to assign a carbon intensity to all other Rystad assets.

B.1 Upstream CI

B.1.1 Flaring data

Upstream emissions exhibit substantial heterogeneity across assets, and the major contributor to this heterogeneity is the flaring-to-oil ratio (FOR). Flaring corresponds to the combustion on-site of co-produced gas[14], and represents more than 150 billion cubic meters worldwide in 2024 (World Bank, 2025). To compute precise estimates of upstream CI, I use satellite data on flares to derive field-specific flaring-to-oil ratios.

Satellite flares observations and imputed emissions. Rystad UCube database does not contain information on gas flared volumes at the asset level. This is mainly due to the fact that most oil and gas operators do not publicly report flared volumes. Moreover, under-reporting by fields operators remains an important issue making the use of self-reported flaring less reliable.

To fill this gap, I use geocoded flaring volumes calculated by the Earth Observation

[14]Co-produced gas is often used to power equipments on-site, or sold, if the asset is connected to the local gas market or possess equipments to liquefy and ship it. If not, excess gas is then either reinjected, to increase pressure in the reservoir, flared, or directly vented. Given the higher global warming potential of methane, flaring is preferable to venting, and venting is banned in most countries except for rare emergencies. Gas reinjection is often costly, making flaring an easier decision in many oil fields.

Group (EOG) at the Payne Institute for Public Policy, within the Colorado School of Mines. The EOG analyzes satellite data from the Visible Infrared Imaging Radiometer Suite (VIIRS), on board of satellites launched in 2012 and 2017, which detects heat from gas flares at night. Flaring volumes are then estimated using the VIIRS Nightfire algorithm (VNF), as detailed by Elvidge et al. (2015) and further refined by Zhizhin et al. (2021). I use the 141,841 observations with positive flared volumes between 2012 and 2024 to assign them to specific assets from Rystad's database.

Rystad assets geographic information. In order to assign flaring observations to Rystad assets, I need geographic information on these assets. I complement data from the UCube with two sets of geocoded data obtained from the same company, one in 2022 and the other in 2025.

First, geocoded boundaries were shared by Rystad in 2022 for 41,762 assets, among which 16,801 were generic (i.e. a 500m radius circle). Since the universe of Rystad assets has evolved between 2022 and 2025, 1,916 generic shapes were added for assets missing in the shapefile, based on their geocoded centroid.

Secondly, data was obtained on the location of wells completed between 2000 and 2025. This dataset contains the latitude and longitude of 1,733,698 producer, injector or wildcat wells, from 39,494 assets. In some regions, the location of wells correspond to the asset centroid, meaning that Rystad probably does not have its precise location. In total, 857,900 unique locations exist in the dataset. Crucially, wells locations are precisely defined for onshore US and Canada assets, where all assets boundaries are generic, making this dataset a useful complement to the first one.

Matching satellite observations to Rystad assets. The assignment of satellite observations to Rystad assets relies on both assets' eligibility (the existence or not of surface wells, excluding for example subsea tiebacks and not yet completed wells) and geographic distances to existing wells and assets' boundaries. For each observation, distances to all completed wells and assets in a 20km radius were computed. A score reflecting the probability of a detected flare to originate from each nearby asset is then computed. For observations with existing wells in a 2km radius, only information on wells was used to compute this score. For other observations, distances to both assets and wells were incorporated in the score. When scores were too close to allow to decide between assets, flared volumes were assigned based on each eligible asset's score and production capacity. Ghosh et al. (2025) provides more detailed information on the 5-stage matching procedure used to assign flaring observations to Rystad assets.

Overall, all observations with at least one eligible well or asset in a 20km radius are matched, leaving 7,261 observations unmatched (corresponding to 4.7% of total flared volumes between 2012 and 2024). Among matched observations, around two thirds were assigned to a unique asset (70.4% of matched flared volumes) and the rest split between multiple assets.

Computing flaring-to-oil ratios (FOR). Once the matching procedure is done, total flared volumes over the 2012-2024 period are computed and divided by total oil production to obtain asset-specific FOR. For not-yet-active assets over this period, an average FOR is assigned. Two categories of assets producing over the period were also attributed averages: assets with at least 75% of years in the sample with positive production outages, mainly due to wars in Lybia and Syria (109 assets), and assets which production started less than 2 years before the end of the period (911 assets). The average FOR is computed at the lowest scale possible, for asset of similar size if possible[15], and always by oil type. Averages are capped at 12,000 scf/bbl, corresponding to twice as more gas flared than oil produced in barrel of oil equivalent. Table 4 reports the number and type of average FORs assigned to Rystad assets, and the corresponding oil reserves.

FOR origin	Number of assets	Share of oil reserves in 2016
From satellite matching	20,571	64.3%
Area x Oil Type x Field Size	5,133	4.2%
Province x Oil Type x Field Size	7,350	7.3%
Province x Oil Type	10,497	8.9%
Country x Oil Type x Field Size	6,662	3.3%
Country x Oil Type	4,250	4.8%
Region x Oil Type x Field Size	2,948	3.4%
Region x Oil Type	3,136	3.1%
World x Oil Type x Field Size	946	0.9%
World x Oil Type	2	0.0%

Table 4: Origin of FOR values.

Table 5 shows the binned distribution of computed FOR. Most assets display relatively low levels of flaring with around a third of them having a FOR equal to 0 (representing around 22% of oil reserves). In these assets, the 0 value can be explained by either complete use of co-produced gas or particularly low gas-to-oil ratio, but it can also be due to flared volumes so low (or so intermittent) that it falls below the satellite detection threshold. Among assets producing mainly oil (assets with gas-to-oil ratios below 10,200 scf/bbl), the reserves

[15]Assets are split in 10 categories depending on their reserve size, with thresholds at 1, 3, 10, 30, 100, 300, 1000, 3000 and 10000 MMboe.

weighted average FOR is 104.4 scf/bbl, and the median is 13.4 scf/bbl.

FOR (kscf/bbl)	Assets count	Reserves (Gbbl)
0	11,651	345.6
0–0.1	10,820	748.7
0.1–1	10,974	435.4
1–10	1,069	21.0
>10	291	2.5

Table 5: Distribution of flaring-to-oil ratios.

B.1.2 OPGEE model

The Oil Production Greenhouse Gas Emissions Estimator (OPGEE) is an engineering-based model developed by the Oil-Climate Index team (Gordon et al., 2015)[16]. It measures GHG emissions from extraction, processing and transport of crude oil in grams of CO2 equivalent per megajoule of oil produced.

Data inputs and matching with Rystad assets. Running the model requires a wide range of inputs for each field, from extraction methods to gas composition. When information is missing, the model computes smart default values (deduced from other variables when possible). I recover publicly available input data from Masnadi et al. (2018) on 958 oil fields worldwide, and match it to Rystad assets manually. Among the 958 original fields, 737 were matched to 1,136 assets in the Rystad database. Most unmatched fields correspond to marginal heavily-depleted oil fields in the US. The difference between the number of OPGEE fields and matched Rystad assets comes from differing level of aggregation. In many instances, a unique OPGEE field is split into multiple assets in Rystad data. In rare occasions, multiple OPGEE fields were matched to a unique Rystad asset.

After completing the matching procedure, I complement the 737 fields dataset with information on 6 variables from Rystad data in order to improve the accuracy of my estimates. These variables are: production level, field age, gas-to-oil ratio, flaring-to-oil ratio, number of producer wells and number of injection wells. Since Rystad databases do not contain information on the quantity of gas reinjected, used on site, or vented, I assume the amount of gas produced to be equal to the sum of flared and sold gas in order to compute gas-to-oil ratios. Finally, extraction methods reported in Masnadi et al. (2018) input data were adjusted if necessary to match information retrieved from Rystad[17].

[16]Latest versions of the model are available at: https://eao.stanford.edu/research-project/opgee-oil-production-greenhouse-gas-emissions-estimator

[17]This is for example the case for water reinjection. If Rystad data reports existing injection wells, at

Computing CI. The resulting input dataset was fed into the OPGEE model to export results. Two distinct carbon emissions accounting methods are available within the model: energy-based allocation (EBA) and carbon product displacement (CPD). In the EBA approach CI are defined such that $CI_{EBA} = \frac{TotalEmissions_{field}}{Prod_{oil}+Prod_{gas}}$, with $Prod_{oil}$ and $Prod_{gas}$ in MJ. This means that total emissions from the oil field are computed, and allocated to oil and gas sold on the market according to their energy content. In the CPD approach, the definition becomes $CI_{CPD} = \frac{TotalEmissions_{field}-Prod_{gas}.\overline{CI}_{gas}}{Prod_{oil}}$. Co-produced gas gives rise to carbon credits corresponding to the energy content of gas produced multiplied by $\overline{CI}_{gas}$, a worldwide average carbon intensity of gas production computed from the CA-GREET model. The latter approach can lead to negative oil CI, which is why the retained approach was the EBA one.

Figure 23 displays CI computed using the model. The production-weighted average carbon intensity is equal to 59 kgCO2eq/bbl (indicated by the dashed red line on the figure). The variance among fields is considerable, with 10% and 90% weighted deciles ranging from 35 kgCO2eq/bbl to 113 kgCO2eq/bbl. Three key features influencing CI are indicated on the graph:

- **FOR**: as expected, large CI values often originate in large quantities of co-produced gas flared. The influence of FOR on CI largely depends on the assumptions made on destruction removal efficiency (DRE). This value measures to the efficiency of flares, i.e. the share of gas effectively combusted and thus transformed into CO2. Since the global warming potential of methane is 28 times larger than the one of carbon dioxide, assumptions on DRE greatly affects computed CI. I keep the default value of 96.1% retained by the model to compute baseline CI, and provide in sensitivity analysis results using values of 91.1% and 98% to cover the range of estimates found in the literature[18](Plant et al., 2022; Jackson and Smith, 2014; Gvakharia et al., 2017; Zavala-Araiza et al., 2021).

- **Steam injection**: Heavier oils (API below 20) tend to be thicker than lighter ones, and often require the injection of steam in order to lower its viscosity and allow its extraction. This is captured by the steam-to-oil ratio (SOR), corresponding to the number of barrels of vaporized water injected to extract one barrel of oil. High SOR values imply larger quantities of liquids extracted, and therefore larger energy consumed to extract oil.

least water or gas injection must be indicated in OPGEE input data. When this was not the case, water reinjection was assumed to happen, since it more common in oil fields than gas reinjection.

[18]While playing a much smaller role than FOR, the quantity of co-produced gas (GOR) can still influence extraction emissions through fugitive emissions of methane related to its processing and transport on site.

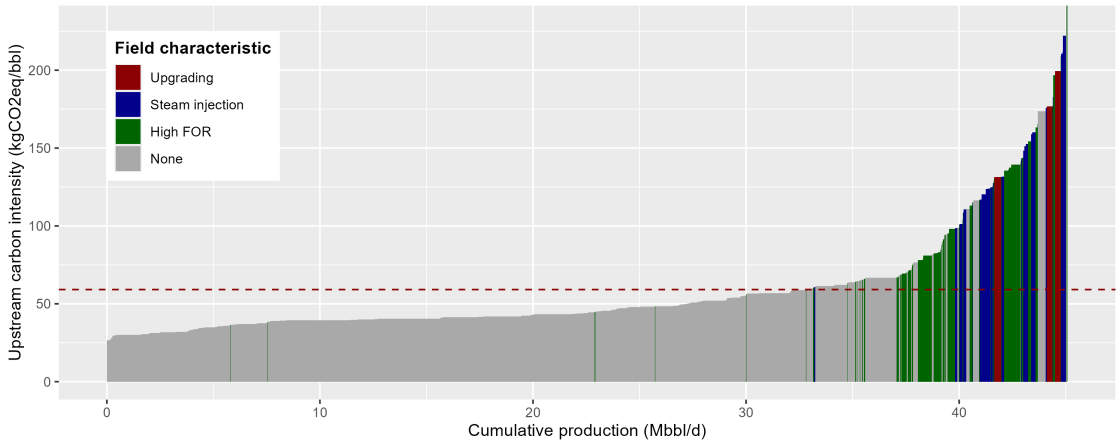

Figure 23: OPGEE-computed CI.

Notes: The figure depicts CI computed in the OPGEE model, with cumulative production in 2015 on the x-axis. The horizontal dashed red line indicates the production-weighted average.

- **Upgrading into synthetic oil**: some particularly viscous oil (bitumen or extra heavy) are extracted without necessitating reducing its viscosity first (for instance using steam injection). This is mainly the case for Canadian oil sands, which are often directly extracted in open air mines. However, they are often too viscous to flow through pipelines and be handled by existing refineries. These oils are therefore "upgraded" into synthetic crude on site. This requires using chemical treatments and/or distillation in order to break long hydrocarbon chains into shorter ones, thus reducing viscosity. These processes are often very energy intensive, translating into large CI for fields using it (Brandt and Farrell, 2007). Produced synthetic crude is however easier to refine, translating into relatively lower midstream CI[19].

B.1.3 Estimation model

Model and sample restrictions. Once the 737 OPGEE-computed CI are attributed to the corresponding 1,136 Rystad assets, I fit a regression model to explain across-field variations. Following the logic of the EBA approach, the dependent variable in the model thus consists of the field carbon intensity divided by its "energy denominator" $ED_f = \frac{Prod_{oil,f}}{Prod_{oil,f}+Prod_{gas,f}}$. The retained model is:

[19]Midstream CI for synthetic crudes are on average 30% lower when compared to other extra heavy or bitumen oils, and still 8% lower when compared to all types of oil.

$$CI_f^{OPGEE} * ED_f = \sum_0^{10} \beta_i^C ShareOilType_f + \beta_{11} Offshore_f + \beta_{12} SOR_f + \beta_{13} GOR_f + \beta_{14} FOR_f + \epsilon_f, \quad (4)$$

where f denotes a Rystad asset, $ShareOilType_f$ the share of each oil type in its production[20], $Offshore_f$ a dummy indicating if the field is offshore, SOR_f its steam-to-oil ratio (in bbl of water/bbl of oil), GOR_f its gas-to-oil ratio (in scf/bbl) and FOR_f its flaring-to-oil ratio (in scf/bbl).

Oil types variables, together with the offshore dummy, capture various aspects of the extraction methods and facility characteristics on site. This is especially relevant for synthetic crude, where upgrading processes such as distillation or chemical treatments significantly increase CI. SOR, GOR and FOR have a more straightforward and linear impact on carbon emissions, and explain most of the remaining variance. SOR data used for the estimation is the one used as input in the OPGEE model since precise values are not available in Rystad data.

The sample used in the regressions was restricted to fields with water-to-oil ratio below 40 bbl of water/bbl of oil, and gas-to-oil ratio below 40,000 scf/bbl. The first restriction excludes fields almost fully depleted, characterized by exponentially increasing WOR[21]. The second one is motivated by the fact that my focus is on oil producing fields. Oil fields are defined by Rystad as reservoirs with GOR inferior to 10,200 scf/bbl. I exclude fields with GOR higher than four times this threshold in order to keep all fields producing non-negligible oil quantities. Altogether, these two restrictions exclude 20 Rystad assets, making the final sample 1,116 assets (associated to 719 OPGEE fields). Weights were added to the regressions in order to keep each unique OPGEE-computed values equivalent. If a unique OPGEE value was matched to multiple Rystad assets, weights were computed based on oil reserves size and sum up to 1.

Regression results. Table 6 shows coefficient estimates from the regression for two sets of CI: the first corresponds to my baseline estimates, using a 100-years horizon for global warming potential, and the second one use a 20-years horizon (associated with higher global warming potential for methane as compared to carbon dioxide, due to its shorter lifetime). R-squared values are close to 1, and predicted R-squared are also quite close to 1, indicating

[20] An asset produces a unique type of crude oil, among 8 possible (synthetic, bitumen, extra heavy, heavy (15-19), heavy (20-22), regular, light and sour). It can however produce simultaneously condensate and/or natural gas liquids (NGL), which are sold as oil, similar to light crude.

[21] I exclude these fields because Rystad data does not contain information on this ratio. The threshold of 40bbl/bbl is particularly high and only excludes 7 fields from the sample.

that the model's predictive power is relatively high.

As expected, synthetic crude production is associated with large CI, and coefficients for heavier crudes tend to be larger than the ones of lighter crudes. While coefficients associated with bitumen and extra heavy crudes are not significantly higher than others, production of these two types often implies the use of steam injection, which means that most of the variation is captured by the SOR variable. The negative coefficient associated with the offshore dummy reflects the absence of land-use transformation, source of GHG emissions. The three remaining variables explain most of the variance. Steam injection greatly increases emissions (the average steam-to-oil ratio is 3 bbl of water/bbl of oil). While having a relatively low magnitude, GOR still plays an important role in explaining the variance in carbon intensity across fields. Fields producing gassy oils, even when selling their co-produced gas, are subject to higher quantities of fugitive emissions. Finally, FOR plays the most crucial role in explaining the variance across fields, as noted by Masnadi et al. (2018).

Regression using 20-years horizon CI as dependent variable display similar results. By construction, reducing the time horizon can only increase CI: while carbon dioxide emissions are not affected, methane emissions lead to higher global warming potential. Consequently, variables associated at least partially with methane emissions see their associated coefficients increase in large proportions. This is particularly visible when comparing the SOR, GOR and FOR variables. While steam injection does not entail methane emissions, gas production leads to fugitive methane emissions and flaring does not fully convert methane into carbon dioxide (depending on the retained DRE value as explained above). The large increase in the coefficients associated to GOR and FOR, along with the SOR coefficient remaining unchanged is thus consistent with this mechanism.

Sensitivity analysis. Figure 24 highlights the strong correlation between OPGEE-computed CI and CI predicted using the model for the same assets sample. To verify the influence of outliers, I re-estimate coefficients of equation 4 excluding each observation one by one. Figure 24 displays estimated coefficients, with the baseline estimate indicated in red. While heavier oils types coefficients can vary due to the low number of assets producing it, results indicate that my estimates are overall robust to changes in the sample. Finally, I test the sensitivity of estimated coefficients to the exclusion of the most influential observations. I compute the Cook's distance of each observation, defined by $CD_i = \sum_{j=1}^{N} \left(\widehat{y}_{\,j} - \widehat{y}_{\,j(i)}\right)^2 / ps^2$, where $\widehat{y}_{\,j}$ and $\widehat{y}_{\,j(i)}$ are the fitted response values from using the full sample and after excluding field i, respectively; s^2 is the mean squared error of the regression model, p the number of coefficients, and N the number of observations. Table 7 displays re-estimated coefficients

Variables	EBA - AR6 100y	EBA - AR6 20y
Oil type: Synthetic	157.096***	192.358***
	(6.514)	(9.943)
Oil type: Bitumen	52.854***	78.442***
	(4.555)	(6.952)
Oil type: Extra Heavy	48.122***	78.368***
	(4.169)	(6.363)
Oil type: Heavy (15-19)	54.778***	96.841***
	(3.73)	(5.693)
Oil type: Heavy (20-22)	58.861***	97.629***
	(3.561)	(5.436)
Oil type: Regular	46.522***	79.708***
	(1.105)	(1.687)
Oil type: Light	46.939***	77.204***
	(1.862)	(2.842)
Oil type: Sour (> 3%)	43.308***	71.654***
	(5.14)	(7.846)
Oil type: Condensate	43.535***	72.446***
	(4.235)	(6.465)
Oil type: NGL	48.05***	83.603***
	(6.641)	(10.137)
Offshore	-6.048***	-8.435***
	(1.273)	(1.943)
SOR (bbl water/bbl oil)	26.317***	27.369***
	(0.983)	(1.501)
GOR (kscf/bbl)	4.601***	8.7***
	(0.154)	(0.236)
FOR (kscf/bbl)	83.908***	146.98***
	(0.533)	(0.813)
Corrected R^2	0.970	0.977
Predicted R^2	0.936	0.943

Table 6: Regression results from upstream emissions estimation

Notes: The table displays coefficients from the regressions using equations 4, with standard errors indicated between parenthesis. R^2 are corrected in order to account for the absence of intercept, as well as predicted R^2.

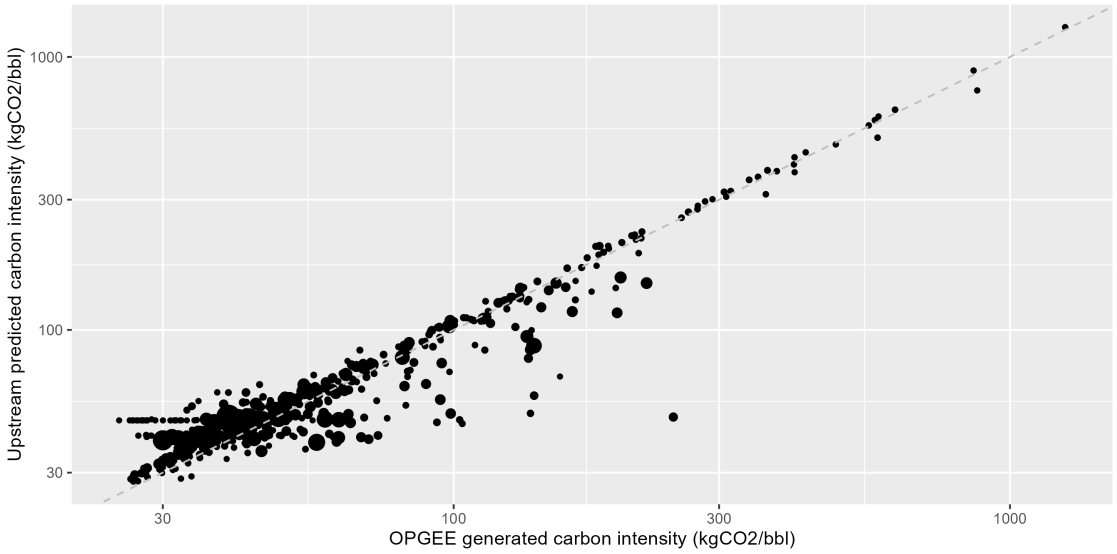

Figure 24: Predicted vs OPGEE-computed upstream CI.

Notes: The figure displays upstream CI computed using the OPGEE model on the x-axis and predicted using the model on the y-axis. Only assets matched with OPGEE fields are included. Both axis are transformed using the decimal logarithm. The size of each point reflects the log volume of oil reserves for each asset.

excluding the 1, 10 and 50 most influential observations. Heavier oil types see their estimate vary slightly, due to the small number of assets producing it in my OPGEE sample, while SOR, GOR and FOR coefficients are particularly stable. Predicted R-squared stay high, even when removing the 50 most influential observations. Overall, re-estimated coefficients remain consistent with my original estimates using the full sample.

B.1.4 Predicting upstream carbon intensity for all Rystad assets

Using estimated coefficients, predicted upstream CI were computed for each Rystad asset. Data on reserves by oil type, offshore status, and GOR was directly available in Rystad databases.Asset-specific FOR were either derived from satellite observations or predicted for not-yet-producing assets, as described in the Flaring data section. Values for other needed variables were assigned as described below:

- SOR: on top of assets matched to OPGEE fields, assets using steam injection were identified and assigned a SOR value of 3, in line with both the default value in the OPGEE model and the average SOR value in OPGEE fields data from Masnadi et al. (2018). All assets containing "Steam", "Thermal" or "SAGD" in their name were identified as using steam injection. In addition, the list of Enhanced Oil Recovery projects made available by the EIA was used to manually identify remaining steam injecting assets (McGlade et al., 2018).

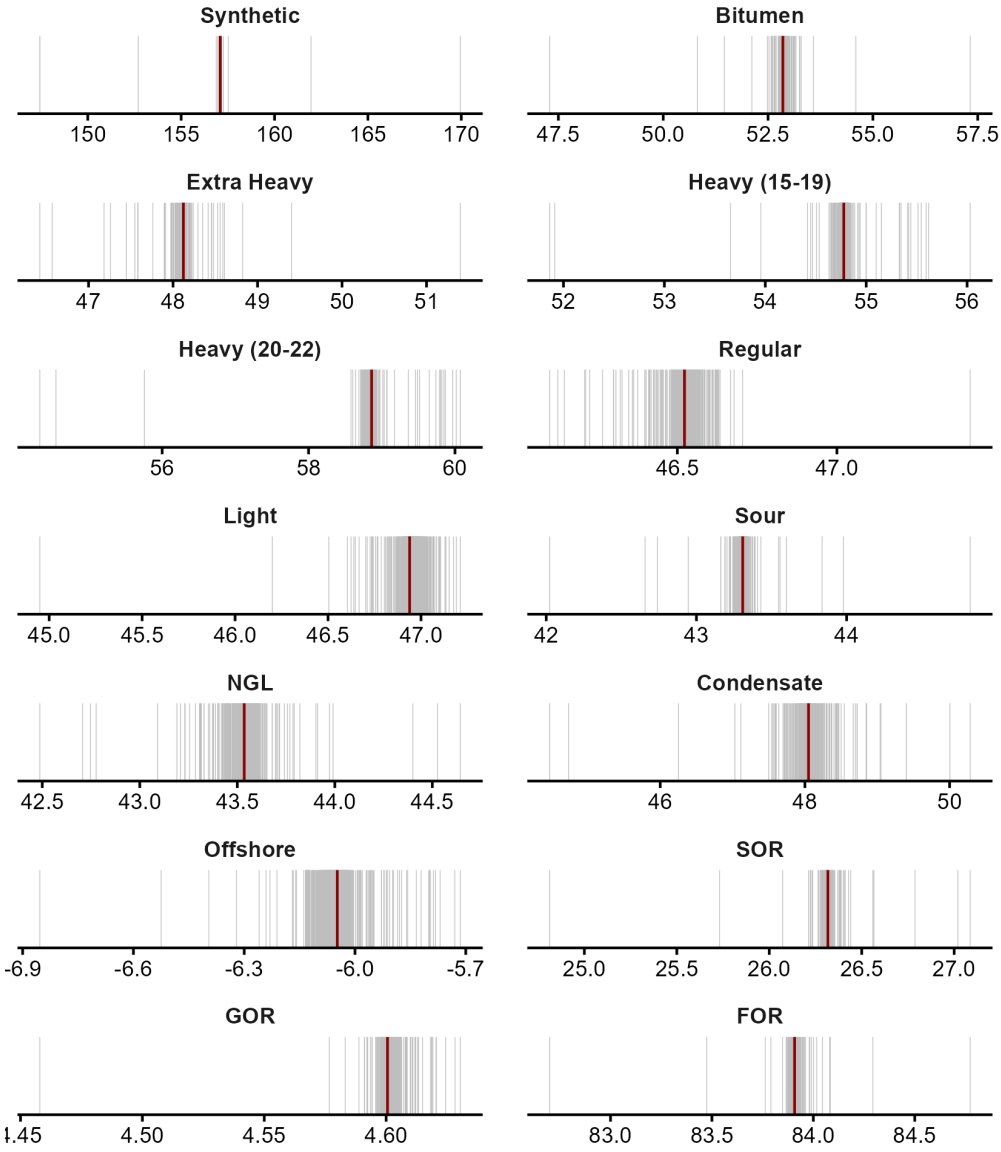

Figure 25: Upstream CI estimation: Estimates of the coefficients on the explanatory variables removing the OPGEE fields one-by-one.

Notes: The panels show the estimated coefficients on each of the explanatory variables in equation (4) removing observations one-by-one. Each grey line corresponds to a different estimation with the same number of observations: $N(OPGEE_{sample})-1$. The thick red lines show the original estimates obtained with the full sample of OPGEE publicly-available fields (OPGEE sample).

	Synthetic	Bitumen	Extra Heavy	Heavy 15–19	Heavy 20–22
0	157.1*** (6.51)	52.85*** (4.55)	48.12*** (4.17)	54.78*** (3.73)	58.86*** (3.56)
1	157.12*** (6.29)	53.14*** (4.40)	48.48*** (4.02)	55.42*** (3.60)	59.49*** (3.44)
10	133.06*** (6.51)	56.86*** (3.70)	51.35*** (3.37)	53.23*** (3.07)	50.40*** (2.99)
50	154.73*** (14.33)	51.56*** (2.88)	47.81*** (2.50)	49.14*** (2.40)	49.60*** (2.64)
	Regular	Light	Sour (>3%)	Condensate	NGL
0	46.52*** (1.11)	46.94*** (1.86)	43.31*** (5.14)	43.54*** (4.24)	48.05*** (6.64)
1	47.42*** (1.07)	46.87*** (1.80)	43.98*** (4.96)	44.40*** (4.09)	49.04*** (6.41)
10	47.62*** (0.90)	45.37*** (1.50)	44.17*** (4.12)	45.48*** (3.41)	53.37*** (5.37)
50	44.59*** (0.65)	41.86*** (1.07)	44.73*** (3.05)	43.83*** (2.51)	46.00*** (3.87)
	Offshore	SOR	GOR	FOR	Pred. R^2
0	-6.05*** (1.27)	26.32*** (0.98)	4.60*** (0.15)	83.91*** (0.53)	0.934
1	-6.85*** (1.23)	26.22*** (0.95)	4.62*** (0.15)	82.70*** (0.53)	0.931
10	-7.39*** (1.03)	24.97*** (0.81)	4.51*** (0.13)	82.67*** (0.54)	0.918
50	-5.21*** (0.74)	26.02*** (0.67)	4.54*** (0.09)	83.93*** (0.46)	0.910

Table 7: Coefficient stability to influential observations (Cook's distance).

Notes: Each panel shows re-estimated coefficients using equation 4 after excluding the 0, 1, 10, and 50 most influential observations. Cook's distances were computed to determine the influence of each observation. The last panel indicates predicted R-squared for each of the re-estimated models.

- ED_f: energy denominators for each asset were approximated by $ED_f = 1 + \frac{GOR_f - FOR_f}{6000}$, using a 6,000MJ average energy content per barrel. This approximation is remarkably close to energy denominators derived from the OPGEE model, with a correlation coefficient of 0.975 within the OPGEE matched assets sample.

When a unique Rystad asset was matched with a unique OPGEE field, the carbon intensity computed using the OPGEE model was kept. For all remaining assets, predicted values prevail. Figure 26 shows the 2012-2024 production-weighted distribution of CI for both samples, with weighted means depicted by the vertical lines. Although averages are remarkably similar, the predictive model does not seem to fully replicate the variance across OPGEE-computed values. The general shape of the distribution remains coherent with OPGEE outputs.

B.2 Midstream-downstream CI

Midstream emissions correspond to emissions from refining crude oil, while downstream emissions include emissions from both transport to end-users and final use (often combustion) of the different oil products obtained from refining. Refining crude oil yields multiple

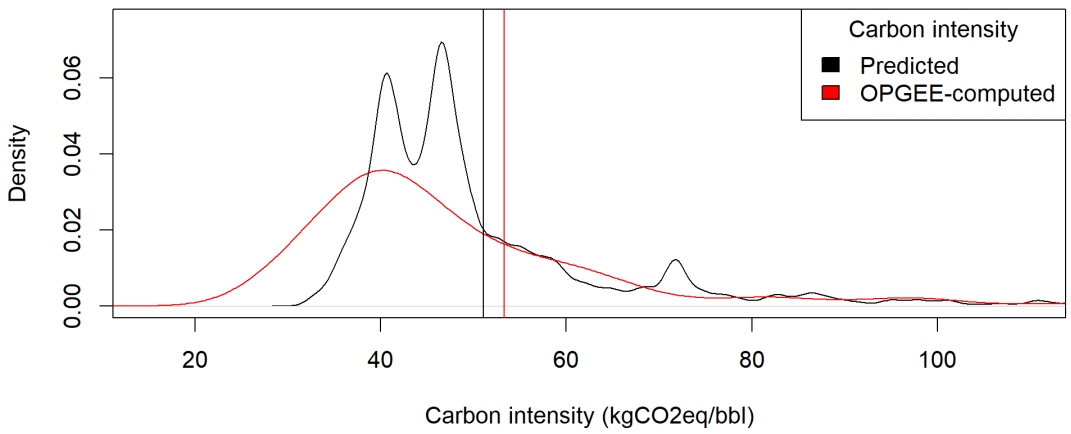

Figure 26: Distribution of CI: Predicted vs OPGEE-computed.

Notes: The figures shows in red the distribution of baseline CI (100-years horizon, EBA) for the OPGEE sample, i.e. assets matched with OPGEE fields and for which the OPGEE value was kept, and in black predicted CI for all remaining assets. Densities are weighted by oil production volumes over the 2012-2024 period. Vertical lines indicate weighted averages for both samples.

oil products depending on its characteristics and refinery configurations. Midstream and downstream emissions are typically computed at the crude stream level, which corresponds to the actual crude stream arriving to the refinery inlet. In some instances, a crude stream originates in a single oil field, while the majority of crude streams consists of multiple fields outputs blended together. For instance, the Brent crude used as reference in oil pricing is a blend of crude oils produced in 13 separate assets in the North Sea.

B.2.1 PRELIM and OPEM models

Midstream and downstream CI can be computed using the Petroleum Refinery Life Cycle Inventory Model (PRELIM) and Oil Products Emissions Module (OPEM) respectively[22]. The versions used in this study to compute CI are PRELIM 1.6 and OPEM 1.1. The PRELIM model requires detailed information on crude assays, such as API gravity, sulphur content, and various boiling points. This data is typically not made available by producers, and publicly available data usually does not follow the same format and is therefore unusable in the models. Models authors thus included in the model data on 657 unique crude assays, that we match to Rystad assets they originate from.

[22]Information and most recent versions of the models can be found at: https://ociplus.rmi.org/methodology.

Running the model. The PRELIM model allows for a large array of refinery configurations, designed to handle different types of oils and aimed at maximizing the production of specific oil products. Since I do not have sufficient information to match Rystad assets to the refineries treating their oil production, I follow the assumptions of the model on refinery configurations. Among the three main configurations used worldwide, the model selects the adapted configuration based on crude oil characteristics in terms of API gravity and sulphur content. All heavy crudes (API gravity below 22) are assigned the "deep conversion" configuration, entailing additional treatments aimed at reducing the length of carbon chains yielded. Light crudes with low sulphur levels (API gravity above 32 and sulphur content below 0.5%) are assigned the "hydroskimming" configuration, associated with less treatment stages and in turn less emissions. All other crudes are refined using the "medium conversion" configuration. The importance of refinery configuration is illustrated in figure 27, depicting PRELIM-computed midstream CI. Intensive treatment of crude in the case of deep conversion is associated with higher emissions, while hydroskimming units yield much smaller emissions.

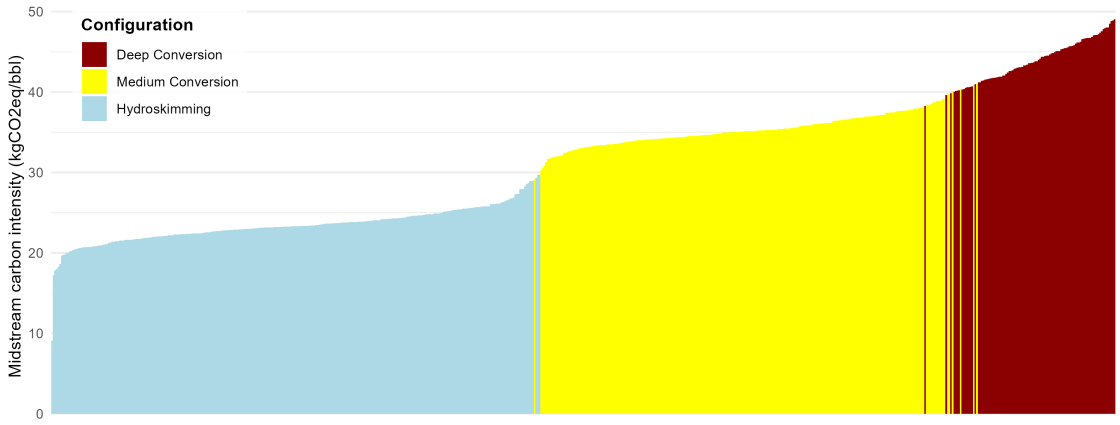

Figure 27: PRELIM-computed midstream CI

Once midstream emissions are computed, PRELIM output sheets are imported to the OPEM model in order to estimate downstream CI of each crude stream. Emissions from transport and final combustion (for oil products used as fuel) are computed, giving me for a given barrel of crude oil its final products yield and associated emissions. Final combustion represents the larger share of emissions along oil lifecycle, and displays smaller variability across crudes.

B.2.2 Matching PRELIM crude assays to Rystad assets

Crude streams available in the PRELIM model are matched to Rystad assets using data made available by Rystad on crude streams and associated assets. The matching is performed using fuzzy matching methods on crudes names. After cleaning crudes names from both sources, Jaro-Winkler distances were calculated and capped at 0.25. Matches were then manually verified, leaving me with 493 crude assays from PRELIM (from the original 657) matched to 392 crude streams from Rystad[23]. Additionally, 25 unmatched PRELIM assets were manually matched to 41 Rystad assets with no specified crude streams. Overall, 518 PRELIM crude assays were linked to 8,386 Rystad assets, representing 64.4% of world production in 2024.

A unique crude stream often consists of a blend of crude oils from multiple assets, but in certain cases (mainly in Canada), production in a unique asset is split between different crudes. For these assets, midstream and downstream CI from PRELIM/OPEM were averaged weighting by the share of production being allocated to each crude stream.

B.2.3 Estimation models

Model and sample restrictions. Using PRELIM output, I fit a regression model to explain the variance in midstream and downstream CI separately. In both cases, I use the two most influential variables, API gravity and sulphur content, which are also available at the asset level in Rystad data. Refinery configuration plays the most important role in explaining variations. I follow the assumptions made in the PRELIM model and assign "deep conversion", "hydroskimming" and "medium conversion" based on the same rules on API gravity and sulphur content. Both variables (in log form for sulphur content) are subsequently introduced in quadratic form to account for as much of the residual variation as possible. Higher-order terms are iteratively dropped to avoid overfitting based on predicted R-squared. The final models used to explain midstream and downstream carbon emissions are as follow:

$$\begin{aligned} CI_f^{PRELIM} &= \sum_1^3 \beta_i^C RefineryConfigf + \beta_4 API_f + \beta_5 API_f^2 + \beta_6 API_f^3 \\ &+ \beta_7 log(Sulphur) + \beta_8 log(Sulphur)^2 + \epsilon_f \end{aligned} \tag{5}$$

[23]The difference in number comes from unique crudes streams being split into different crude assays in PRELIM differing by the data source. For instance, the "Hibernia" crude in Canada is split into 3 assays, corresponding to information on this specific crude obtained from ExxonMobil, Statoil and Chevron.

$$CI_f^{OPEM} = \sum_1^3 \beta_i^C RefineryConfigf + \beta_4 API_f + \beta_5 API_f^2 + \beta_6 API_f^3 + \beta_7 API_f^4 + \beta_8 log(Sulphur) + \beta_9 log(Sulphur)^2 + \beta_{10} log(Sulphur)^3 + \epsilon_f \quad (6)$$

Regression results. Table 8 shows results from both estimations. All coefficients are significant and predicted R-squared are 0.95 and 0.88 respectively. As expected, refinery configuration explains a large share of variations across crudes, especially in the case of midstream emissions. Adding quadratically API gravity and log sulphur content improves the variance explained marginally, as the relatively low within R-squared indicates. On the contrary, API gravity and sulphur content play a much more important role in explaining the variance in downstream emissions.

Variables	Midstream	Downstream
Configuration: Deep conversion	56.552***	530.269***
	(1.374)	(17.341)
Configuration: Hydroskimming	40.093***	587.056***
	(1.608)	(17.764)
Configuration: Medium conversion	51.361***	515.626***
	(1.604)	(18.037)
API	-0.933***	-16.122***
	(0.124)	(2.159)
API^2	0.016***	0.781***
	(0.003)	(0.096)
API^3	-9.62e-05***	-0.017***
	(2.74e-05)	(0.002)
API^4		1.11e-04***
		(1.16e-05)
log(Sulphur)	-0.988***	-7.784***
	(0.103)	(0.832)
$log(Sulphur)^2$	-0.135***	-2.404***
	(0.025)	(0.474)
$log(Sulphur)^3$		-0.19***
		(0.072)
Corrected R^2	0.948	0.887
Predicted R^2	0.947	0.883
Within R^2	0.300	0.875

Table 8: Regressions results from midstream and downstream CI estimations.

Notes: The table displays coefficients from the regressions using equations 5 and 6, with standard errors indicated between parenthesis. R^2 are corrected in order to account for the absence of intercept, as well as predicted R^2. Within R^2 are computed using fixed-effects regressions to capture the effect of configuration dummies.

Sensitivity analysis. The distribution of midstream and downstream CI are compared to predicted values using the model in figure 28. Midstream values appear grouped according to each refinery configuration, illustrating the crucial role played by the dummies in the regression results. While the inclusion of API gravity and sulphur content in the model allows to explain a large share of within group variation for medium and deep conversion, it does not appear to be the case for crudes using the hydroskimming configuration (bottom-left of panel A), explaining the low within R-squared value. Downstream CI are more evenly distributed across the graph, and predicted values remain close to OPEM-computed ones.

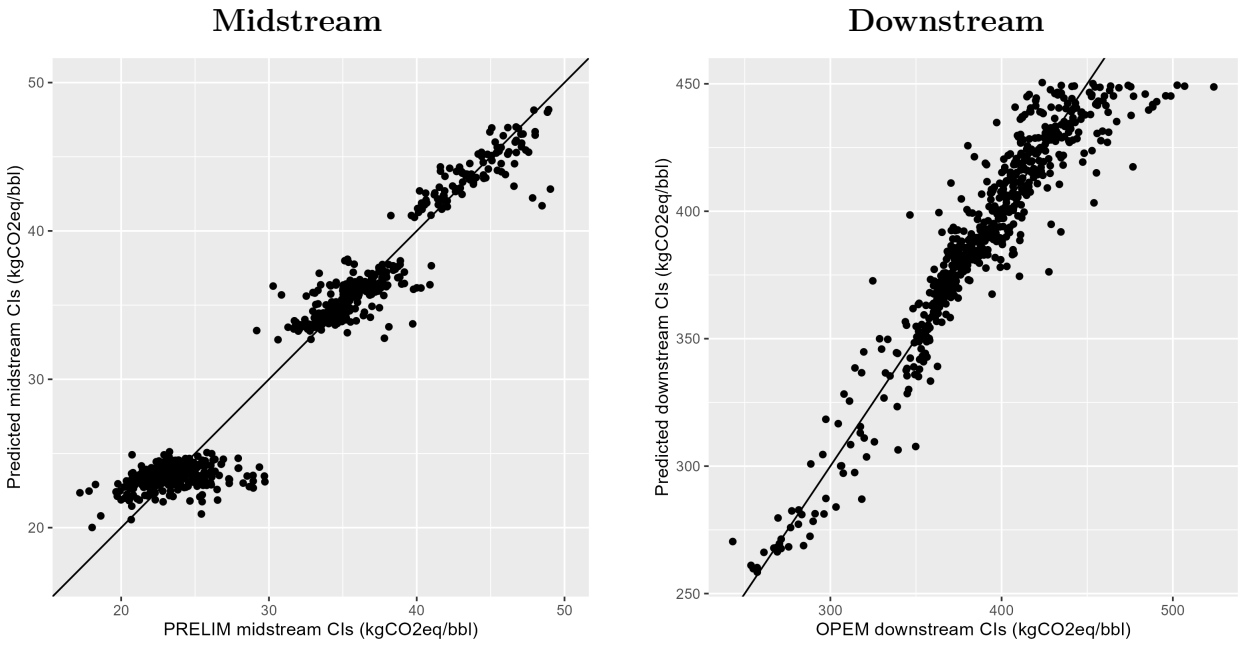

Figure 28: Comparison between PRELIM/OPEM CI and predicted ones.

Contrary to upstream CI, midstream and downstream CI computed using PRELIM and OPEM models do not contain visible outliers. The risk of overfitting particular values is thus low. Tables 9 and 10 reports re-estimated coefficients excluding most-influential observations for midstream and downstream CI respectively. In both cases, estimates remain significant and stable enough to exclude the risk of overfitting.

B.2.4 Predicting midstream/downstream CI

Rystad assets matched to PRELIM crudes are assigned midstream and downstream CIs computed using the PRELIM and OPEM models. These assets represent 59% of oil reserves in 2016. For remaining assets, I use the coefficients estimated in both regressions to

	Deep Conversion	Hydroskimming	Medium Conversion	API
0	56.55*** (1.37)	40.09*** (1.61)	51.36*** (1.60)	-0.93*** (0.12)
1	56.38*** (1.30)	40.05*** (1.52)	51.22*** (1.51)	-0.91*** (0.12)
10	56.69*** (1.23)	40.58*** (1.44)	51.73*** (1.44)	-0.93*** (0.11)
50	56.56*** (1.05)	40.80*** (1.22)	51.91*** (1.21)	-0.89*** (0.10)

	API^2	API^3	log(Sulphur)	$log(Sulphur)^2$
0	0.0159***	-9.62e−05***	-0.99*** (0.10)	-0.13*** (0.03)
1	0.0152***	-9.03e−05***	-1.00*** (0.10)	-0.13*** (0.02)
10	0.0153***	-8.77e−05***	-1.10*** (0.09)	-0.18*** (0.02)
50	0.0126***	-5.57e−05***	-1.18*** (0.08)	-0.22*** (0.02)

Table 9: Midstream CI estimation: coefficient stability to influential observations (Cook's distance).

Notes: Each panel reports coefficients re-estimated after excluding the 0, 1, 10, and 50 most influential observations according to Cook's distance. Standard errors are reported in parentheses.

	Deep Conversion	Hydroskimming	Medium Conversion	API	API^2
0	530.27*** (17.34)	587.06*** (17.76)	515.63*** (18.04)	-16.12*** (2.16)	0.78*** (0.10)
1	543.82*** (18.09)	599.01*** (18.32)	527.93*** (18.62)	-18.39*** (2.33)	0.90*** (0.11)
10	544.71*** (16.65)	600.12*** (16.83)	528.63*** (17.12)	-18.46*** (2.15)	0.91*** (0.10)
50	555.93*** (14.05)	609.59*** (14.18)	540.02*** (14.41)	-19.59*** (1.83)	0.94*** (0.09)

	API^3	API^4	log(Sulphur)	$log(Sulphur)^2$	$log(Sulphur)^3$
0	-1.66e−02***	1.11e−04***	-7.78*** (0.83)	-2.40*** (0.47)	-0.19*** (0.07)
1	-1.92e−02***	1.30e−04***	-7.81*** (0.83)	-2.24*** (0.48)	-0.16** (0.07)
10	-1.93e−02***	1.31e−04***	-7.60*** (0.76)	-2.40*** (0.45)	-0.15** (0.07)
50	-1.95e−02***	1.30e−04***	-7.31*** (0.60)	-2.74*** (0.36)	-0.21*** (0.06)

Table 10: Coefficient stability to influential observations (Cook's distance).

Notes: Each panel reports coefficients re-estimated after excluding the 0, 1, 10, and 50 most influential observations according to Cook's distance. Standard errors are reported in parentheses.

predict their CI, at the crude level in the case of crude streams unmatched to PRELIM crude assays (17.3% of reserves) and at the asset level if the crude stream is unspecified (11.7%).

Several assets, representing 12% of world reserves, lack information on either API gravity or sulphur content. For these assets, reserve-weighted averages from similar oils were computed, according to the following order: reservoir, sub-basin, similar API within the same country, same main oil type within the same country, and finally same main oil type. The vast majority of assets are attributed values computed within the same reservoir or sub-basin. Geographically close assets having a greater probability to be blended together before refining, averages computed from assets nearby is preferable. Table 11 sums up the methods and granularity levels of midstream and downstream CI predictions.

Method	Granularity level	Share of 2016 reserves
PRELIM/OPEM	Crude stream	58.9%
Predicted	Crude stream	17.3%
	Asset-specific	11.7%
Average	Reservoir	5.4%
	Sub-basin	5.7%
	API category × country	0.7%
	Oil type × country	0.2%
	Oil type	$< 0.1\%$

Table 11: Data sources and aggregation levels for midstream and downstream carbon intensity values.

References

Brandt, A. R. and Farrell, A. E. (2007). Scraping the bottom of the barrel: greenhouse gas emission consequences of a transition to low-quality and synthetic petroleum resources. *Climatic Change*, 84(3):241–263.

Elvidge, C. D., Zhizhin, M., Baugh, K., Hsu, F.-C., and Ghosh, T. (2015). Methods for global survey of natural gas flaring from visible infrared imaging radiometer suite data. *Energies*, 9(1):14.

Ghosh, S., Jean, L., and Saha, P. (2025). The last good owner effect: Ownership structure and environmental outcomes in oil and gas.

Gordon, D., Brandt, A. R., Bergerson, J., and Koomey, J. (2015). *Know your oil: creating a global oil-climate index*. Carnegie Endowment for International Peace Washington, DC.

Gvakharia, A., Kort, E. A., Brandt, A., Peischl, J., Ryerson, T. B., Schwarz, J. P., Smith, M. L., and Sweeney, C. (2017). Methane, black carbon, and ethane emissions from natural gas flares in the bakken shale, north dakota. *Environmental Science & Technology*, 51(9):5317–5325.

Jackson, P. M. and Smith, L. K. (2014). Exploring the undulating plateau: the future of global oil supply. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2006).

Masnadi, M. S., El-Houjeiri, H. M., Schunack, D., Li, Y., Englander, J. G., Badahdah, A., Monfort, J.-C., Anderson, J. E., Wallington, T. J., Bergerson, J. A., et al. (2018). Global carbon intensity of crude oil production. *Science*, 361(6405):851–853.

McGlade, C., Sondak, G., and Han, M. (2018). Whatever happened to enhanced oil recovery?

Plant, G., Kort, E. A., Brandt, A. R., Chen, Y., Fordice, G., Gorchov Negron, A. M., Schwietzke, S., Smith, M., and Zavala-Araiza, D. (2022). Inefficient and unlit natural gas flares both emit large quantities of methane. *Science*, 377(6614):1566–1571.

Welsby, D., Price, J., Pye, S., and Ekins, P. (2021). Unextractable fossil fuels in a 1.5 °C world. *Nature*, 597(7875):230–234.

World Bank (2025). Global gas flaring tracker report, july 2025. Technical report, World Bank.

Zavala-Araiza, D., Omara, M., Gautam, R., Smith, M. L., Pandey, S., Aben, I., Almanza-Veloz, V., Conley, S., Houweling, S., Kort, E. A., et al. (2021). A tale of two regions: methane emissions from oil and gas production in offshore/onshore mexico. *Environmental Research Letters*, 16(2):024019.

Zhizhin, M., Matveev, A., Ghosh, T., Hsu, F.-C., Howells, M., and Elvidge, C. (2021). Measuring gas flaring in russia with multispectral viirs nightfire. *Remote Sensing*, 13(16):3078.