Second-best carbon taxation in a differentiated oligopoly

Diego S. Cardoso*

August 2026

Abstract

This paper examines second-best carbon taxation in the US domestic aviation sector, where a corrective tax interacts with two pre-existing distortions: oligopoly market power and non-environmental taxes. I document novel empirical evidence of the large effects of these imperfections on the efficiency of carbon taxes. Combining accounting aggregates, sufficient statistics, and structural modeling in three nested levels of aggregation, I estimate social abatement costs, calculate the optimal carbon tax under both distortions, and investigate the effects of introducing a revenue-neutral carbon tax. I find that the current social gross marginal abatement cost is \$196–\$243/ton CO_2, implying that for a social cost of carbon (SCC) below this range, any positive carbon tax reduces social welfare. The three levels agree near the margin, but coarser implementations overstate average abatement costs by nearly 80% for tax changes of realistic magnitude. Even under a high SCC of \$300/ton CO_2, the optimal carbon tax is only \$114/ton CO_2, thus much lower than the standard Pigouvian benchmark. Moreover, substituting the existing sales tax with a revenue-neutral carbon tax would not yield a double dividend: removing the sales tax deadweight loss raises welfare, but firms adjust markups to absorb the reallocated tax burden, leaving aggregate emissions virtually unchanged.

*University of Illinois Urbana-Champaign, Department of Agricultural and Consumer Economics, 1301 West Gregory Drive, Urbana, IL, 61801, USA. Email: dcardoso@illinois.edu. I am grateful to Volodymyr Bilotkach, Julien Daubanes, Todd Gerarden, Nathaniel Hendren, Jonathan Hughes, Catherine Kling, Shanjun Li, Yongjoon Park, Ivan Rudik, and Ruozi Song for their helpful comments and suggestions. I also thank participants at the AERE, ASSA, CREEA, and EEA conferences, the Northeast Workshop on Energy Policy and Environmental Economics, and seminar participants at Cornell University, Duke University, Hamilton College, Indiana University, Institute for Policy Integrity, Michigan State University, Norwegian School of Economics, Syracuse University, University of Connecticut, University of Geneva, University of Illinois Urbana-Champaign, and ZEW.

1 Introduction

"Levy a tax equal to the marginal external cost" is a foundational policy prescription in the economic analysis of externalities. This type of tax, known as Pigouvian taxation, makes agents internalize the external costs, maximizes social welfare, and leads markets to an efficient equilibrium in the absence of other distortions. As such, Pigouvian taxes and equivalent market-based instruments have enjoyed broad support among economists as a policy to mitigate environmental damages.

The optimality of a Pigouvian tax, however, hinges on the assumption that the target externality is the only market imperfection (Buchanan, 1969). Nevertheless, industries that are disproportionate sources of pollution also tend to be highly concentrated—the cases of cement, concrete, steel, oil, and electricity, for example. This systematic presence of market power in major polluting sectors has drawn continuous attention from Industrial Organization, and Environmental and Public Economics research (Kellogg & Reguant, 2021). The convergence of interests across these fields is not coincidental: quantifying how market power interacts with environmental policy is critical to designing optimal corrective taxes.

This paper examines second-best carbon taxation in a well-known oligopoly: the US domestic aviation sector. Commercial aviation is a notoriously concentrated sector that has proven challenging for climate policy. A rich literature has documented that airlines exercise market power as a result of the oligopolistic nature of the sector (e.g., Borenstein, 1989; Ciliberto & Williams, 2010). Adding to market power distortions, air travel is subject to non-Pigouvian distortionary taxes. The sector also generates climate-related externalities, accounting for approximately 3% of global greenhouse gas emissions and 5% of the radiative forcing leading to climate change (Lee et al., 2009). With limited regulation, carbon emissions from aviation are projected to continue growing (Owen et al., 2010) and may account for as much as 22% of net greenhouse gas emissions by 2050 (European Parliament, 2015). Despite efforts to curb emissions from international aviation, the scope of policies has been limited, and large domestic air travel markets have not been addressed. Most notably, the US—the largest aviation market—has only recently devised climate action plans, albeit primarily focused on broad, long-term goals targeting technological advancements (FAA, 2021).

Using data from the US Department of Transportation, this paper implements a common welfare framework at three levels of aggregation, each requiring progressively more detailed data. At the coarsest level, a sector-wide calculation recovers the key distortion measures from publicly reported financial aggregates, with no estimation required. At the intermediate level, a sufficient statistics approach combines market-level aggregates with an estimated demand elasticity to capture heterogeneity across markets. At the finest level, a structural model of demand and competition recovers product-level equilibrium responses, which is required for evaluating non-marginal tax changes. Based on these nested implementations, I (i) derive marginal and non-marginal welfare costs of emission abatement via carbon taxation; (ii) calculate the second-best carbon tax taking market power and existing taxes as given; and (iii) examine whether a revenue-neutral carbon tax in place of existing sales taxes can generate welfare gains while reducing emissions.

This paper finds that any positive carbon tax would decrease social welfare if the social cost of carbon[1] (SCC) is equal to \$50 per metric ton of CO_2 (henceforth, tCO_2). Under a higher SCC of \$300/tCO_2—an upper-bound estimate from Rennert et al. (2022)—the optimal carbon tax is approximately \$114/tCO_2. Therefore, the optimal tax level is about a third of the standard Pigouvian prescription. The striking difference between the optimal tax and the marginal damage represented by the SCC can be mapped to existing market imperfections.

Estimates across the three approaches indicate that abating one tCO_2 by increasing carbon taxes would lead to a loss of \$139–185 in private surplus due to markups and \$50–54 due to sales tax distortions. These substantial distortions accentuate the loss of private surplus when carbon taxes increase, and further analysis shows that these losses are evenly split between reductions in consumer and producer surplus. As a result, if emission reductions are achieved exclusively through demand suppression via taxation, the social gross marginal abatement cost is between \$196 and \$243/tCO_2 at the baseline. The agreement across implementations with very different requirements is itself informative, as even a back-of-the-envelope calculation requiring no estimation delivers an abatement cost near the center of this range. The levels diverge, however, once tax changes are large: only the structural model captures how markups adjust and demand reallocates away from carbon-intensive flights, and a comparison exercise shows that linear extrapolation of the aggregated approaches understates emission responses under a carbon tax comparable to recent fuel price swings. The comparison across levels also shows that the sufficient statistics implementation overstates the markup wedge by about a quarter, mainly because accounting markups are proportional to distance—and hence to carbon intensity—by construction.

Welfare theory suggests that replacing a tax based on prices for one based on the externality

[1]The social cost of carbon indicates the present discounted value of the stream of climate damages from an additional ton of CO_2 emissions.

could increase the efficiency of taxation. The current sales tax, set at 7.5% of the fare, partially substitutes the role of a carbon tax by increasing average ticket prices and reducing demand and emissions. However, sales taxes can be inefficient instruments because fares are not necessarily linked to emissions. In this case, a revenue-neutral carbon tax could result in the so-called double dividend—lower emissions and increased welfare by replacing a more distortive tax. I investigate this alternative in counterfactual analyses that implement a tax reform replacing the current sales tax with a carbon tax.

In line with theory, I find that a revenue-neutral carbon tax of about \$62/tCO_2 would increase aggregate welfare. Surprisingly, however, the double dividend does not materialize because this tax substitution leads to insignificant reductions in aggregate emissions. Hence, net social welfare gains follow primarily from increases in private surplus. This result is driven by the presence of market power: when a carbon tax shifts the tax burden to more polluting flights, firms can respond by reducing markups to retain market share. Lower markups and the removal of the sales tax deadweight increase private surplus. However, markup reductions also undo part of the effect on prices intended by taxing carbon. As a result, emissions among the most polluting flights decrease less than in perfect competition, and these reductions are commensurate with the total emission increase from the less polluting options.

These findings highlight one of the main challenges for climate policy in the sector: with limited economically-viable abatement technologies in the short run, abating carbon emissions via demand reduction instruments has a high cost. Moreover, with existing distortions, attempts to reform tax instruments in order to achieve emission targets might backfire. Over a longer horizon, technological advancements and alternative fuels are more likely to reduce the cost of abatement and make emission targets more feasible. In the past few years, several initiatives have been set in place to incentivize the development and expansion of sustainable aviation fuels.[2] These policies, combined with the results presented here, suggest that a comprehensive climate policy for the aviation sector should consider a mix of demand-side and supply-side instruments.

This paper makes three main contributions. First, it provides the first analysis of how market

[2] Lohawala, Toman, and Joiner (2026) present a detailed discussion of the potential of sustainable aviation fuels (SAFs) as well as current and future policies. Despite showing great promise to steer aviation towards net-zero goals, SAFs will have limited impact in the short and medium run. Current production pathways cost of 4 to 8 times more than conventional jet fuels. Moreover, scaling up production might lead to increased competition for limited feedstock supply capacity. Even if the optimistic targets set by the US SAF Grand Challenge of increasing annual production from about 16 million in 2022 to 3 billion gallons by 2030 were met, SAFs would supply only about 11% of the projected domestic fuel demand (DOE, 2021).

power and environmental externalities jointly influence the economic efficiency of environmental taxation in the airline industry. Since the seminal work of Buchanan (1969) showing how the standard Pigouvian tax can lead to welfare loss in monopoly, a strand of the literature has advanced theoretical models to characterize the interplay of environmental externalities and market power (Requate, 2006). Empirical work on this interaction has emerged more recently, typically focusing on a single industry; it has shown, for instance, that existing market power can lead to emission reductions in electricity (Mansur, 2007), environmental regulations can exacerbate market power in cement (Ryan, 2012), and market power and incomplete cost pass-through reduce the efficiency of carbon taxes in coal (Preonas, 2024); incomplete pass-through has also been documented in several other carbon-intensive sectors (e.g., Lade & Bushnell, 2019; Ganapati et al., 2020; Muehlegger & Sweeney, 2022). For aviation, previous studies have predicted the effects of hypothetical carbon policies on prices and demand—carbon taxes (Brueckner & Abreu, 2017; Pagoni & Psaraki-Kalouptsidi, 2016), cap-and-trade programs (Winchester et al., 2013), and jet fuel tax increases (Fukui & Miyoshi, 2017)—but have overlooked the economic efficiency of policy instruments and the role of non-environmental market imperfections. This paper provides a comprehensive welfare analysis of carbon taxation for this sector, quantifying the efficiency loss of policy instruments under imperfect competition.

Second, this paper accounts for pre-existing taxation in the analysis of market power and externalities. In doing so, it contributes to the literature on second-best externality taxation, much of which has examined the conditions for a double dividend in general equilibrium settings with distortionary input taxes (e.g., Bovenberg & Mooij, 1994; Bovenberg & Goulder, 1996; Cremer et al., 1998; Goulder et al., 1999). Under pre-existing taxation, optimal environmental taxes are frequently smaller than the marginal damage (Bovenberg & Mooij, 1994; Parry, 1995), distortionary taxes increase the cost of environmental policy (Goulder et al., 1999), and monopoly power can intensify this effect (Fullerton & Metcalf, 2002). Beyond environmental applications, recent work in empirical public economics has quantified how market power reshapes commodity taxation. Pre-tax prices can be strategic substitutes for the tax rate, so that firms capture most of the revenue a tax reform was intended to raise (Miravete, Seim, & Thurk, 2018), and a uniform rate levied on differentiated products underprices some of them, implicitly subsidizing the consumers and firms concentrated in those products (Miravete, Seim, & Thurk, 2020). This paper adds empirical evidence on the triple interaction between market power, pre-existing taxes, and environmental externalities in a differentiated oligopoly; here, pre-existing sales tax acts as a fiscal externality channel for the carbon tax (Hendren & Sprung-Keyser, 2020), and a revenue-neutral substitution between the two instruments can raise welfare without significantly reducing emissions.

Third, this paper demonstrates how sufficient statistics and structural approaches can be combined—and nested—to assess marginal and non-marginal effects of environmental policies in imperfect markets. The sufficient statistics approach evaluates marginal welfare changes from a policy based on a small set of estimable parameters (Chetty, 2009); previous environmental applications include the efficiency costs of price instruments (Jacobsen, Knittel, Sallee, & Van Benthem, 2020) and distributional effects in imperfect markets (Genakos, Grey, & Ritz, 2020). Determining the optimal tax, however, involves evaluating non-marginal changes and requires additional structural assumptions (Kleven, 2021). This paper develops these insights into a three-level empirical framework in which the same welfare effects quantification is implemented with sector aggregates, market-level sufficient statistics, and a product-level structural model. The nesting clarifies what each layer of data and structure offers: the coarser levels provide transparent benchmarks and local consistency checks for the structural model, while the structural model extends the analysis to the non-marginal changes that policy evaluation requires.

The remainder of this paper is organized as follows. Section 2 summarizes key characteristics of the US aviation sector and the challenges they present for climate policy. Section 3 introduces the welfare framework used to characterize market distortions. Section 4 describes the data. Section 5 presents the empirical strategy at three levels of aggregation, and Section 6 reports the estimation results and assesses the fit of the structural model. The welfare analyses of carbon taxation are presented in Section 7. Section 8 offers concluding remarks.

2 US aviation and climate change

The US has the largest domestic commercial aviation market in the world. It accounts for approximately 30% of all passengers carried and 45% of domestic passenger-miles served on domestic flights (IATA, 2019). In growth, the US domestic aviation market is second only to China (IATA, 2019). After the deregulation of commercial aviation in the late 1970s, the sector has experienced tremendous expansion in service, growing from around 250 billion revenue passenger-miles (RPM) in the early 1980s to over a trillion RPMs per year in the late 2010s. This section highlights three aspects of the aviation industry relevant to this paper: market structure, existing taxes, and carbon emissions.

Sector structure. Since deregulation, the industry has seen changes in its players, with various rounds of entry, bankruptcy, and consolidation.[3] Airlines can be broadly organized into two groups (Belobaba et al., 2015). One group includes the *legacy airlines*, alluding to the fact that these firms have operated since the pre-deregulation era. This group currently includes American, Delta, and United Airlines. Some of the distinguishing features of these players are their extensive service networks with large hubs, more rigid cost structures with higher levels of unionization, and bundled higher-quality services (such as meals and in-flight amenities). The other group is formed by *low-cost carriers* (LCCs), which follow the "no-frills" business model successfully implemented by Southwest Airlines. Examples of other airlines in this category are Spirit and Allegiant. As the name suggests, LCCs focus on running cost-efficient operations. This involves, for example, flying point-to-point services from smaller airports instead of maintaining large hubs, keeping a high aircraft utilization rate, and unbundling passenger services by charging extra fees for baggage, food, beverages, and other amenities.

In practice, the legacy vs. LCC categorization is more of a conceptual construction than an accurate description of how airlines operate. The financial success of LCCs has led legacy carriers to adopt some of the LCC practices. On the opposite end, the quest for diversification has also led some LCCs, such as JetBlue, to invest in higher service quality. Furthermore, a third group of airlines can be categorized as regional carriers; these players are either small, independent companies that run on limited networks or carriers that operate in partnership with larger airlines to provide connections from hubs to smaller airports—under the brand name of United Express or American Eagle, for example (Belobaba et al., 2015).

With the small number of airlines and high fixed and entry costs, the aviation industry largely operates as an oligopoly. Extensive literature has documented evidence of market power in the US aviation industry. Though a review of this literature is beyond the scope of this brief description, prior findings present some common themes. For example, studies have found that the existence of a hub premium is a source of market power (Borenstein, 1989, 1991; Lee & Luengo-Prado, 2005; Lederman, 2007, 2008; Berry & Jia, 2010). Other mechanisms generating and maintaining market power are tacit collusion (Evans & Kessides, 1994; Ciliberto & Williams, 2014; Ciliberto et al., 2019; Aryal et al., 2022), entry deterrence (Ciliberto & Williams, 2010; Aguirregabiria & Ho, 2012; Ciliberto & Zhang, 2017), and mergers and consolidation (Kim & Singal, 1993). In the

[3] Borenstein and Rose (2014) present a comprehensive overview of the US aviation industry, including its history, trends, and unique challenges.

opposite direction, increased competition from LCCs, a trend initially attributed to the "Southwest effect," has been found to reduce prices and markups (Morrison, 2001; Goolsbee & Syverson, 2008; Brueckner et al., 2013).

Taxes and fees. In the US, commercial flights are subject to a sales tax, a fuel tax, and various fees. The sales tax corresponds to the US Federal Excise Ticket Tax, set at 7.5% of the base fare. This tax is dedicated to the Airport and Airway Trust Fund (AATF), which most notably helps to fund the Federal Aviation Administration (FAA, 2020). In addition, jet fuel used for commercial aviation is subject to a federal tax of 4.3 cents/gallon, also appropriated by the AATF, plus a 0.1 cents/gallon fee, appropriated by the Leaking Underground Storage Tank Trust Fund. There are also three fees for domestic flights in the US: (i) the Federal Security Surcharge, at \$11.20 per domestic round-trip itinerary; (ii) the Federal Flight Segment Tax, which charges \$4.20 per domestic segment; and (iii) Passenger Facility Charges, costing on average \$4.50 per departing airport. Notably, the model proposed in this paper captures these taxes and fees, as they have implications for pricing behavior.

Emissions. Aviation accounts for 2–3% of global greenhouse gas annual emissions (Owen et al., 2010) and is one of the sectors with the fastest growth in emissions. Between 1990 and 2016, greenhouse gas emissions from aviation grew by 98% (FCCC, 2018). These emissions are projected to grow by 200–360% in the first half of the 21st century (Owen et al., 2010). With the fast expansion of air travel and limited abatement alternatives, the sector lags behind other industries in decarbonization. As a result, aviation may globally account for up to 22% of net greenhouse gas emissions by 2050 (European Parliament, 2015).

Aviation's total contribution to climate change is larger than its share of CO_2 emissions, accounting for as much as 5% of the radiative forcing leading to global warming (Lee et al., 2009). Jet fuel burn releases other components that affect heat transfer in the atmosphere, including water vapor, nitrous oxides, soot, and contrails. Though some components favor atmospheric cooling, such as aerosols and methane reduction due to nitrous oxides, the net effect of jet fuel burn produces warming above the individual contribution of CO_2 (Lee et al., 2009).

3 Welfare analysis framework

The welfare analysis of taxation under imperfect competition rests on well-developed foundations (see Requate (2006) for a survey of environmental applications). Most directly related to this paper, Weyl and Fabinger (2013) characterize the incidence of unit taxes in general models of imperfect competition, showing how pass-through summarizes the division of tax burdens, and Adachi and Fabinger (2022) extend this framework to the joint incidence of unit and *ad valorem* taxes with multiproduct firms. Both papers evaluate welfare as the sum of consumer surplus from quasi-linear utility, producer surplus as operating profits, and tax revenue, and both approach taxation as a classical revenue-raising problem, focused on incidence and the marginal value of public funds.

This section builds on that framework to address a different question: the welfare cost of abating an externality through taxation when markets already feature market power and distortionary taxes. Two modifications are required by this question and by the application. First, the welfare aggregate incorporates environmental damages, which allows the framework to characterize the social gross marginal abatement cost (SGMAC) and to decompose the gap between the second-best carbon tax and marginal damages into markup and sales tax wedges. Second, the tax system combines a uniform *ad valorem* sales tax with an emissions tax that is effectively heterogeneous across products, since tax payments are proportional to each product's emission intensity. None of these components is novel in isolation, as they are standard concepts in the public and environmental economics literatures. The value of assembling them in a single framework is that it makes explicit which objects must be measured to operationalize the welfare formulas.

3.1 Setting

Markets and products are defined following common practices the aviation literature (Berry, Carnall, & Spiller, 2006; Aguirregabiria & Ho, 2012). Time is discrete, and each period t represents a quarter. A *location* is a city or metropolitan area with one or more airports, and a *market* m is a directional pair of locations.[4] A *product* is the combination of a route operated by airline i in market m at period t.[5] To simplify notation, let k index the K_m total products available in a given

[4] Pagoni and Psaraki-Kalouptsidi (2016) adopt a similar approach but define locations as clusters of airports within a radius, while other studies define markets as directional airport pairs (e.g., Borenstein, 1989; Ciliberto & Tamer, 2009; Berry & Jia, 2010). Defining markets over metro areas allows the model to capture competition between flights departing from airports in close proximity.

[5] In this paper, routes are defined as round-trip itinerary with up to one connection each way. Flights with disjoint segments or more than one connection each way are not considered; in the data used for estimation, they correspond

market.

Each product k has an emission intensity $e_k = h \times f_k$, where h is the carbon intensity of jet fuel ($h = 0.0134$ tCO_2/gallon)[6] and f_k is fuel burn per available seat. Airlines supply seats while consumers purchase trips, and the load factor linking these two units is held fixed throughout, a restriction discussed in Section 5.3 and documented in Appendix D. Assuming no economically-viable abatement technology exists in the short run, emission intensities are fixed and all abatement is achieved through consumption reduction.[7] The carbon tax τ, measured in dollars per tCO_2, is implemented as a volumetric tax on jet fuel. Such taxation approach is a natural choice in this setting because jet fuel is a homogeneous, single-use commodity whose combustion maps directly into emissions, and a federal jet fuel tax already exists, thus lowering the institutional requirements for implementation.[8] Since tax payments per unit equal τe_k, the carbon tax operates as a product-specific unit tax whose level varies with emission intensity.

Beyond the carbon tax, each flight is subject to a uniform *ad valorem* sales tax r and a product-specific fee ι_k, so the ticket price paid by consumers is $p_k = (1 + r)\,\tilde{p}_k + \iota_k$, where $\tilde{p}_k$ is the pre-tax price received by the airline. Fees ι_k represent infrastructure costs, such as airport and security services provided by entities other than the airlines. These fees are not taxes, but they create a wedge between the price paid and the amount received by airlines and must therefore be included in the model (see Section 2 for an overview of these charges).

Social welfare in this setting comprises four components: consumer surplus, firms' operating profits, tax revenue,[9] and environmental damages.

Consumer surplus. As in Weyl and Fabinger (2013) and Adachi and Fabinger (2022), consumer surplus derives from quasi-linear utility, so policy-induced income effects are absent. There are N_m identical consumers with utility $U_m\left(x_0, x_1, \ldots, x_{K_m}\right) = \alpha_m x_0 + \sum_k u_k\left(x_k\right)$, where x_0 is a composite outside good (i.e., not flying) with unit price that generates no externality, and α_m determines the

to less than 3% of all domestic enplanements.

[6]Each gallon of jet fuel burned emits an average of 9.57 kg CO2 (EIA, 2016). To account for other fuel burn components contributing to climate change, these emissions are converted to CO_2-equivalent using a 1.4 conversion factor from (Azar & Johansson, 2012).

[7]This assumption is relaxed in Appendix C.1, which extends the model to allow for input substitution.

[8]Given the fixed carbon intensity h, an emissions tax of τ dollars per tCO_2 corresponds to a volumetric tax of τh dollars per gallon. Isomorphic alternatives, such as tradable permits or ticket taxes indexed to emissions, are not separately analyzed.

[9]I make no assumptions about how firm profits are distributed to individuals or how tax revenues are recycled, for which reason these elements are kept in separate accounts.

marginal utility of income. Utility maximization defines demands $x_k^* = x_k(\boldsymbol{P}_m, y)$, where $\boldsymbol{P}_m$ is the price vector and y is income. Letting $q_k(\tau)$ and $p_k(\tau)$ represent the equilibrium quantity and price under emissions tax τ, the standard aggregate money-metric consumer surplus in market m can be written as $CS_m(\tau) = \alpha_m^{-1} N_m \sum_k u_k(q_k(\tau)/N_m) - \sum_k p_k(\tau) q_k(\tau) + N_m y$.

Operating profits. There are I_m firms supplying the emission-generating products; the outside option is competitively supplied and can be abstracted from profit considerations. Let c_k denote the constant marginal cost of production excluding carbon tax payments. Operating profits[10] for product k are given by $\pi_k(\tau) = [\tilde{p}_k(\tau) - c_k - \tau e_k] q_k(\tau)$. The term in square brackets defines the *operating markup*, $\mu_k \equiv \tilde{p}_k - c_k - \tau e_k$, the wedge between the pre-tax price and marginal cost inclusive of carbon tax payments. Then, the total operating profits in a market are denoted by $\Pi_m(\tau) = \sum_k \pi_k(\tau)$.

Tax revenue. Total tax revenue combines receipts from the emissions tax and the sales tax: $T_m(\tau) = \tau \sum_k q_k(\tau) e_k + r \sum_k q_k(\tau) \tilde{p}_k(\tau)$.

This expression highlights the potential for fiscal externalities due to the interaction of two taxes (Hendren & Sprung-Keyser, 2020): an increase in the carbon tax induces behavioral responses in quantities that can lower the revenue raised from sales taxes.

Environmental damage. Each unit of emissions produces a constant environmental damage ϕ, measured here exclusively for carbon-equivalent emissions and corresponding to the social cost of carbon. Environmental damages in market m are $\Phi_m(\tau) = \phi \sum_k q_k(\tau) e_k$.

Social welfare. Aggregate net social welfare is the unweighted sum of the four components: $W_m(\tau) = CS_m(\tau) + \Pi_m(\tau) + T_m(\tau) - \Phi_m(\tau)$.

[10] These are, in fact, variable operating profits. Throughout the paper, airlines' route networks and fleets are held fixed, which limits the analysis to short-run effects: entry and exit in response to a tax change lie outside the framework, for which reason I also omit fixed costs. Over a longer horizon, airlines adjust networks based on dynamic and strategic considerations (Belobaba et al., 2015), and modeling those decisions requires a substantially more complex framework. Appendix C.2 bounds the consequences of assuming no tax-induced entry and exit.

3.2 Marginal welfare effects

Let $W \equiv \sum_m W_m$ represent aggregate social welfare in a sector with multiple markets. Differentiating W with respect to τ at the baseline tax level yields

$$\frac{dW}{d\tau} = \sum_k \left[r\tilde{p}_k + \mu_k + (\tau - \phi)\, e_k\right] \frac{dq_k}{d\tau}, \tag{1}$$

where the summation over k now includes all products in the sector. The terms within square brackets show how all three market imperfections affect welfare. The sales tax and markup terms are analogous to the usual Harberger triangle terms (Kleven, 2021): for these, welfare losses equal the mechanical variation in tax revenue and operating profit. The third term captures the environmental externality and its correction mechanism.

Two properties of this expression organize the empirical work that follows. First, the terms involving price derivatives cancel across welfare components, because a marginal change in the pre-tax price transfers $(1 + r)\, q_k$ away from consumers and delivers exactly that amount to firms and the government combined.[11] This cancellation uses only consumer optimization; no firm first-order condition enters, and no assumption is required on how markups adjust to the tax. Second, and as a consequence, the framework is agnostic about the origin of the demand responses $\frac{dq_k}{d\tau}$. This agnosticism allows the same expressions to be evaluated under different sets of restrictions, for which reason no particular model of competition is specified at this stage.

3.3 Social marginal abatement cost and second-best carbon tax

The ratio between the marginal change in private surplus with the marginal change in aggregate emissions defines the *social gross marginal abatement cost* (SGMAC, following the terminology in Phaneuf and Requate, 2016):

$$SGMAC\,(\tau) \equiv \tau + \frac{\sum_k \left[r\tilde{p}_k + \mu_k\right] \frac{dq_k}{d\tau}}{\sum_k e_k \frac{dq_k}{d\tau}}. \tag{2}$$

The SGMAC measures the aggregate private welfare loss per tCO_2 abated when emission reductions are induced by raising τ and no abatement technology exists.[12] The second term shows that market

[11] See Appendix A.1 for the derivation. Markup adjustments still matter for welfare, as they operate through the equilibrium quantity responses $\frac{dq_k}{d\tau}$.

[12] Emission intensities e_k are held fixed, so emission reductions are obtained exclusively through reduced consumption. Appendix C.1 derives the SGMAC when production allows substitution toward cleaner inputs. Carbon offsets

imperfections raise abatement costs above the tax itself: each ton abated through demand reduction also destroys the markup and sales-tax revenue associated with the forgone units. The same objects characterize the second-best tax. Setting $\frac{dW}{d\tau} = 0$ in equation (1) yields

$$\tau^* = \phi - \underbrace{\left\{ \frac{\sum_k \mu_k \frac{dq_k}{d\tau}}{\sum_k e_k \frac{dq_k}{d\tau}} \right\}}_{\text{Markup wedge}} - \underbrace{\left\{ \frac{\sum_k r\tilde{p}_k \frac{dq_k}{d\tau}}{\sum_k e_k \frac{dq_k}{d\tau}} \right\}}_{\text{Sales tax wedge}}, \tag{3}$$

the familiar result that non-environmental distortions drive a wedge between the optimal environmental tax and marginal damages (Buchanan, 1969). Throughout the text, I refer to the two components as the *markup* and *sales tax wedges*, measured in $\$/\text{tCO}_2$; standard Pigouvian taxation arises when $r = \mu_k = 0$, and combining (2) and (3) gives $SGMAC\,(\tau^*) = \phi$. The empirical content of this framework lies in magnitude of the wedges relative to the externality.

3.4 Tax substitution and the double dividend

The existence of a distortive sales tax creates the possibility of a double dividend: replacing the sales tax with an externality tax could both decrease emissions and improve welfare. In this case, it is assumed that the current level of tax revenue cannot be reduced, as it funds necessary government expenditures. This restriction implies a change in the sales tax rate (dr) for every change in the emissions tax ($d\tau$) such that total revenue is unchanged, $dT = 0$. Both instruments are set nationally, so the constraint binds on sector aggregates, for which reason the sums below run over all products in the sector. Let $t_k \equiv \tau e_k + r\tilde{p}_k$ denote the tax paid per unit of product k, and let $E \equiv \sum_k q_k e_k$ and $B \equiv \sum_k q_k \tilde{p}_k$ denote aggregate emissions and the sales tax base. Totally differentiating $T = \sum_k q_k t_k$ with respect to either instrument yields

$$\frac{\partial T}{\partial x} = M_x + \sum_k t_k \frac{\partial q_k}{\partial x} + r \sum_k q_k \frac{\partial \tilde{p}_k}{\partial x}, \qquad x \in \{\tau, r\}, \quad M_\tau = E, \quad M_r = B, \tag{4}$$

and revenue neutrality requires $\frac{dr}{d\tau} = -\left(\frac{\partial T}{\partial \tau}\right) / \left(\frac{\partial T}{\partial r}\right)$. The two derivatives differ only in the mechanical term M_x. The behavioral terms capture the impacts on tax bases through quantity adjustment and the movement of the sales tax base induced by pre-tax price responses.[13]

or emission permits connected to sectors with lower abatement costs would induce additional net reductions, thus differing from the analysis of an aviation-only carbon tax.

[13] When markups are invariant to the instruments, the last term in each expression vanishes. Appendix A.2 derives (4) and reports the restricted case, maintained by the sector- and market-level implementations in Section 5, in which

Here, a simplification can provide intuition on the substitution channels. Setting the behavioral terms aside gives the mechanical rate of substitution, $\frac{dr}{d\tau} = -E/B$. Defining the *implied carbon tax* of product k as $\theta_k \equiv r\tilde{p}_k/e_k$, the sales tax paid per tCO_2 emitted, this benchmark equals $-r/\bar{\theta}$, where $\bar{\theta} \equiv rB/E$ is the emissions-weighted mean of the θ_k. Evaluating the mechanical change in per-unit tax payments at this rate yields the condition

$$\left.\frac{dt_k}{d\tau}\right|_{\text{mechanical}} \geq 0 \Longleftrightarrow \theta_k \leq \bar{\theta}.$$

Absent behavioral effects, products whose implied carbon tax lies below the emissions-weighted average pay more after the substitution, and those above it pay less. Since a product's pre-tax price reflects willingness to pay for characteristics rather than its carbon content, θ_k is not exactly proportional to e_k, and the substitution reallocates tax burdens toward carbon-intensive flights by construction.

A complete analysis of tax substitution requires the welfare and emission responses along this path. Imposing $dT = 0$ leaves the marginal welfare expression (1) unchanged in form, with $\frac{dq_k}{d\tau}$ reinterpreted as the total derivative along the revenue-neutral path, as shown in Appendix A.2. The substitution therefore changes the magnitude of the demand responses but not the objects that must be measured. It does not, however, determine the sign of $dE = \sum_k e_k \frac{dq_k}{d\tau}$. Two channels impact the reallocation of burdens and the response of quantities. The first is pass-through, which under differentiated-product competition is neither diagonal nor equal to the mechanical burden change and may even exceed unity (Anderson, de Palma, & Kreider, 2001). The second is substitution, which reallocates demand across products whose emission intensities are unordered. Neither channel can be signed a priori (Weyl & Fabinger, 2013; Adachi & Fabinger, 2022). The environmental dividend depends on the distribution of θ_k relative to e_k and pass-through that preserves the intended reallocation, but neither condition is guaranteed by theory. Therefore, whether the double dividend holds is an empirical question, which Section 7.3 addresses for this sector.

3.5 From welfare formulas to measurement

Equations (1)–(3) are expressed in terms of product-level prices, emission intensities, markups, and demand responses. Prices and emission intensities are directly observed but markups and demand responses are not. Estimating markups requires knowledge of marginal costs, and demand

markups are invariant to both taxes and demand depends only on own price.

responses require the full matrix of own- and cross-price effects, which is generally infeasible to estimate with the data available for this sector. Therefore, operationalizing the framework requires imposing restrictions, and the choice among restrictions must consider the level of detail the data can offer and what the resulting estimates can answer. Section 5 develops three implementations of the same formulas at increasing levels of detail.

4 Data

Data sources. The data set used in this paper combines seven data sources from four providers. US aviation data are sourced from the Bureau of Transportation Statistics (BTS) of the US Department of Transportation (BTS, 2018). I query four BTS databases in this paper. First, the Origin and Destination Survey (DB1B) provides quarterly data on domestic air travel—including origins, destinations, connections, and ticket prices—based on a 10% sample of all tickets. Second, Table T-100 of the Form 41 Traffic Database contains monthly data on air travel operations by segment, aircraft, and airline. From this database, I collect the number of available seats, passengers transported, and ramp-to-ramp time by segment, aircraft model, and airline; I also collect the use share of each aircraft model by segment and airline. Third, I gather data on the number of departures and delays by segment and airline from the On-Time Performance Database. Fourth, I collect aggregate operating revenues and costs by airline and quarter from the Form 41 Financial Database, Schedule P-1.2.

Average fuel burn by aircraft model and stage length (i.e., flight segment distance) are collected from the International Civil Aviation Organization's carbon emissions calculator documentation (ICAO, 2018). Jet fuel prices come from the US Energy Information Administration (EIA, 2018); I collect monthly prices of sales to end users by region. Finally, data on population and personal income per capita at the Metropolitan Statistical Area are collected from the US Bureau of Economic Analysis' Regional Economic Accounts (BEA, 2018).[14]

Sample criteria. The sample used in this analysis includes the four quarters of 2018. The data set covers all 73 cities or metropolitan areas in the contiguous US with at least 50,000 passengers surveyed in 2018, which jointly account for 92% of all domestic traffic. These locations include a total of 98 airports.

[14]Additional details on the construction of the data set are presented in Appendix D.

Following the data reliability criteria adopted in the literature (Berry & Jia, 2010; Aguirregabiria & Ho, 2012; Pagoni & Psaraki-Kalouptsidi, 2016), itineraries in the DB1B data are excluded if any of the following conditions apply: (i) is operated by a non-US airline; (ii) is not a round trip; (iii) has more than one stop in either direction; (iv) has fare credibility issues flagged by the BTS; (v) has extreme fares (below \$50 or above \$3,000); or (vi) has fewer than 3 tickets surveyed in a quarter.

Itineraries selected from the DB1B are assigned to the reporting airline, which is the same as the operating and ticketing airlines in the majority of the cases. Moreover, small and regional carriers with exclusive service for or acquired by another airline are grouped with their controlling airline.

Covariates. The resulting data set has 267,967 observations, each representing one product. These products are offered by 10 airline groups in 5,018 markets. Table 1 presents summary statistics for the covariates used for estimation of model parameters; financial aggregates from Form 41 are included in Table 5, which is also used to compare reported and model-predicted metrics.

Some covariates require additional details. *Shares* are calculated based on *market size*, which is defined as the geometric mean of the origin and destination populations, as in Berry and Jia (2010) and Pagoni and Psaraki-Kalouptsidi (2016). *Passengers* is the number of passengers for a product surveyed in DB1B multiplied by 10 (the survey weight). *Market distance* is the great circle distance between the origin and destination airports. *Connection extra distance* is the travel distance added by having connections; it is calculated as the difference between total travel distance and twice the market distance. *Price* is the quarterly mean, post-tax ticket price for a product. *Departures per week* is assigned to the minimum number of departures across each segment in a route. *Delayed flights* indicates the percent of departures in the previous quarter with delays above 15 minutes. *Destinations from origin* indicates the number of destinations an airline serves from the origin airport. *Total ramp-to-ramp time* sums the time taken in each segment of a route.

Seven covariates are used as excluded instruments for estimation, as explained in section 5. This set of instruments was chosen based on previous studies on the aviation sector (especially Berry et al., 2006; Berry & Jia, 2010; Aguirregabiria & Ho, 2012; Pagoni & Psaraki-Kalouptsidi, 2016). Five of these instruments measure the degree of competition: the number of *airlines in market*, the number of *rivals' products in market*, the percentage of those rivals' products that are nonstop

Table 1: Summary statistics of estimation sample.

Statistic	Mean	St. Dev.	Min	Median	Max
Share (%)	0.01	0.06	0.0002	0.002	2.04
Share within nest (%)	7.37	15.41	0.01	1.41	100.00
Passengers	509.25	2,261.27	30	60	62,640
Market Size ($\times 10^{-6}$)	3.47	2.40	0.43	2.76	16.02
Origin income per capita ($1000s)	59.04	12.42	35.84	56.28	99.42
Price ($)	479.06	157.41	56.97	465.33	2,041.33
Number of stops	1.60	0.64	0	2	2
Market Distance (miles)	1,371.16	640.40	67.13	1,256.22	2,724.08
Connection distance (miles)	273.36	295.94	0.00	180.99	2,992.77
Departures per week	15.85	13.54	0.08	13.00	145.69
Delayed flights (%)	17.28	5.77	0.00	16.88	100.00
Destinations from origin	25.09	19.40	1	18	84
Airlines in market	5.38	1.48	1	5	9
Rivals' products in market	57.34	69.44	0	34	604
Rivals' % of nonstop flights	5.00	9.44	0.00	2.53	100.00
Potential legacy entrants	0.05	0.25	0	0	3
Potential LCC entrants	3.01	1.06	0	3	5
Compl. segment density ($\times 10^{-3}$)	48.69	28.42	0.02	44.25	256.40
Fuel expenditure ($/avail. seat)	97.46	31.77	6.44	94.93	276.49
Ramp-to-ramp time (h)	8.36	2.75	1.50	8.03	18.40
Observations (products)			267,967		
Routes			103,720		
Time periods (quarters)			4		
Airlines			10		
Markets			5,018		
Cities or metropolitan areas			73		
Airports			98		

flights, and the number of legacy and low-cost *potential entrants*. Similar to Goolsbee and Syverson (2008), potential entrants are identified as airlines currently not offering flights in a market but operating in the origin or destination cities. *Complementary segment density* indicates the sum of passengers from other markets who are transported on each segment of a route; this variable is a measure of the scale of operations along a route. Finally, *fuel expenditure* is the sum of fuel consumption per available seat along each segment, multiplied by the respective fuel price.[15]

[15]This covariate accounts for the length of each segment, as well as fuel efficiency and use share of each aircraft model in each segment of a route. Jet fuel prices are assigned based on the departing airport in each segment. See Appendix D for details.

5 Empirical strategy

The welfare characterizations of Section 3 are portable because they leave demand responses $\frac{dq_k}{d\tau}$ unrestricted. This means that three key objects can be evaluated under different sets of restrictions on markups and demand responses: the wedges in equation (3), the SGMAC in equation (2), and the marginal welfare changes in equation (1). This section develops three such sets, ordered from the least to the most data-demanding. At the sector level, all flights form a single composite product and no estimation is required. At the market level, flights within a market form a composite product, and a single estimated elasticity translates observed market weights into welfare effects. At the product level, a specified model of demand and competition recovers the full set of equilibrium objects. The levels are nested, in the sense that relaxing the aggregation restriction of one level yields the next, and each additional layer of structure buys the ability to answer questions that the coarser levels cannot.

Each level carries assumptions that may be innocuous for some quantities and consequential for others, so stating them separately for each implementation is critical when interpreting their different results. The three subsections that follow present each level in turn, together with the measurement or estimation it requires. Section 5.4 then places them side by side and states what each one can and cannot deliver. Estimation results are reported separately in Section 6.

5.1 Level 1: Sector aggregation with reported statistics

At the coarsest level, all domestic flights form a single composite product with sector-average pre-tax price $\bar{\bar{p}}$, operating markup $\bar{\mu}$, and emission intensity $\bar{e}$. A single product leaves no composition margin, so the demand response is a scalar that enters both the numerator and the denominator of the wedge terms and cancels:

$$SGMAC\left(\tau\right) = \tau + \frac{\left(r\bar{\bar{p}} + \bar{\mu}\right)\frac{dQ}{d\tau}}{\bar{e}\frac{dQ}{d\tau}} = \tau + \frac{r\bar{\bar{p}} + \bar{\mu}}{\bar{e}}. \tag{5}$$

The markup and sales tax wedges of equation (3) reduce correspondingly to $\bar{\mu}/\bar{e}$ and $r\bar{\bar{p}}/\bar{e}$. No estimation is required at this level: the wedges are ratios of accounting averages. This is the empirical content of the framework in its most compact form: a back-of-the-envelope calculation available wherever sector-level financial aggregates exist.

The inputs come from the system-wide financial aggregates that airlines report in the Form

41, described in Section 4: operating revenues and operating costs. Here, the unit of output is defined as available seat-miles (ASM), a metric widely used in aviation to quantify aggregate supply. Dividing the difference between aggregate operating revenues and costs by the aggregate ASM gives the sector-average markup per ASM. Emission intensity per ASM is calculated from fuel consumption and the carbon intensity of jet fuel h. Since all three inputs to equation (5) are expressed per ASM, the normalization cancels in the ratios, and the wedges follow from the reported averages together with the sales tax rate r.

Two approximations deserve mention. First, operating revenue includes ancillary sources (baggage, reservation, and cancellation fees) that are not subject to the ticket excise tax, so a sales tax wedge computed from total revenue per ASM slightly overstates the taxed base. Second, sector averages weight the entire domestic network, while the market- and product-level analyses use a sample that skews toward larger markets. Both features are acceptable for the purpose this level serves. However, aggregation at this level imposes a consequential limitation: the metrics overlook heterogeneity of any kind, cannot indicate which markets are most affected, and force quantities and emissions to change in exact proportion, since a single composite product admits no composition effects.

5.2 Level 2: Market aggregation with reported and sufficient statistics

Fully characterizing product-level demand responses in each market requires the full matrix of own- and cross-price effects. Such estimation is infeasible with the data available in this context, as each market features its own set of products and the number of parameters becomes intractable. This level works around this limitation by modeling flights within a market as a composite product, while allowing markets to remain distinct products. Quantities within a market are assumed to move proportionally following a marginal tax change, so that $dq_k = s_{k|g}\, dQ_{mt}$, where $s_{k|g}$ is the observed within-market share of product k. Within-market composition is therefore fixed, while cross-market composition responds to differences in emission intensities, prices, and pass-through. Flights in market m at period t are treated as a composite product $Q_{mt} = \sum_k q_k$, with average price p_{mt}, fuel use f_{mt}, and markup μ_{mt} weighted by within-market shares.

Implementing this level requires markups, pass-through elasticities, and the elasticity of aggregate demand. Product or market-level markups (or equivalently, marginal costs) are not observed, and recovering granular markups typically involves a structural approach based on an equilibrium concept; this is the approach described in Section 5.3. Alternatively, mean markups can be approximated using detailed financial data on revenue and costs. For the aviation sector, such statistics are available at the firm level through the Form 41, including operating revenues, costs, and ASM per carrier. Hence, I use these metrics to calculate carrier-level markups per ASM. The calculation is similar to the sector level, but now assigned at the firm level, so that markups vary across flights from the same airline in proportion to travel distance. This measurement, together with the aggregation restriction above, embeds two assumptions: (i) markups are approximated by the difference between company-level revenues and costs per seat-mile; and (ii) market shares and markups are invariant to small fuel price changes. Assumption (i) is the most consequential, because the correlation between markups and emission intensity across markets partly drives the welfare conclusions; Section 7 uses the structural model's equilibrium markups to assess the direction and magnitude of the bias it introduces. Assumption (ii) is a standard marginal approximation, but here implies that marginal increments in fuel costs are fully passed through to ticket prices; its role, however, is narrower than it appears for the reason given below.

Under assumption (ii), a marginal increase in the carbon tax raises fuel expenditure per available seat by $\frac{dF_{mt}}{d\tau} = hf_{mt} = e_{mt}$ and translates into ticket prices as

$$\frac{dp_{mt}}{d\tau} = \frac{dp_{mt}}{dF_{mt}} \frac{dF_{mt}}{d\tau} = \eta_{mt}\, p_{mt} \frac{e_{mt}}{F_{mt}}, \tag{6}$$

where $F_{mt} = (w_{mt} + \tau h)\, f_{mt}$ is market-average fuel expenditure per available seat, w_{mt} is the jet fuel price per gallon, and $\eta_{mt} \equiv \frac{\partial \ln p_{mt}}{\partial \ln F_{mt}}$ is the pass-through elasticity of fuel costs to ticket prices. Under complete marginal pass-through, this elasticity equals the fuel cost share of the ticket price and can be calculated directly from observed variables as $\eta_{mt} = (1+r)\, F_{mt}/p_{mt}$, so no estimation is involved. Aggregating equation (6) under the proportional-shares restriction yields the response of the market composite,[16]

$$\frac{dQ_{mt}}{d\tau} = \varepsilon\, \eta_{mt} \frac{e_{mt}}{F_{mt}} Q_{mt}, \tag{7}$$

where $\varepsilon = \frac{\partial \ln Q_{mt}}{\partial \ln p_{mt}}$ is the elasticity of market-aggregate demand with respect to the market-average ticket price.[17]

A useful property follows from ε being common across markets. Since this parameter multiplies every market response in equation (7), it cancels in the ratios that define the wedges and the SGMAC, just as the scalar response cancels at the sector level. Cross-market heterogeneity in the

[16] See Appendix A for the derivation of the sufficient statistics and their use in calculating marginal welfare changes.

[17] Since $e_{mt} = hf_{mt}$, the ratio $e_{mt}/F_{mt} = h/\left(w_{mt} + \tau h\right)$ converts the volumetric fuel cost base into emission units.

wedges is therefore captured by observed market-level weights alone (quantities, fuel shares, and emission intensities) and the estimated elasticity is required only to express marginal changes in welfare, quantities, and emissions in levels. The same argument applies to assumption (ii): uniform assumptions on the absolute pass-through rate are innocuous for the wedges, which depend only on observed weights, and matter only for changes expressed in levels.

Estimation and identification. The remaining input is the elasticity of aggregate demand, which I estimate in a reduced-form approach with

$$\ln Q_{mt} = \varepsilon \ln p_{mt} + \beta^{(\varepsilon)} X_{mt} + \gamma_m^{(\varepsilon)} + \gamma_t^{(\varepsilon)} + \nu_{mt}^{(\varepsilon)}, \tag{8}$$

where X_{mt} is a vector of time-varying, market-average product characteristics, $\gamma_m^{(\varepsilon)}$ and $\gamma_t^{(\varepsilon)}$ represent fixed effects for non-directional city pairs and quarters, and $\nu_{mt}^{(\varepsilon)}$ is an idiosyncratic demand shock. These fixed effects capture demand components related to characteristics of the endpoint locations and to seasonality.

As usual in demand estimation, the market price p_{mt} is likely endogenous. To address this bias, I instrument it with the average fuel expenditure per available seat, a cost shifter constructed as in the aviation literature (e.g., Berry et al., 2006; Pagoni & Psaraki-Kalouptsidi, 2016). Identification relies on the assumption that consumer choice does not respond systematically to changes in fuel costs except through the effect of those costs on prices. This instrument combines fuel burn per available seat along the segments of a route with the regional price of jet fuel at the departing airport of each segment. Regional prices differ in level but closely track movements in spot prices (Figure A-7). The critical identifying variation, thus, comes from the different exposures to price shocks through the heterogeneity of fuel burn per available seat, which depends on stage length, aircraft model, and the number of segments a passenger flies (see Figure A-6).[18]

5.3 Level 3: Multiproduct firms with structural modeling

At the finest level, no aggregation is imposed. Products are airline-route-quarter offerings, firms are asymmetric multiproduct suppliers, and both markups and demand responses are equilibrium objects. Recovering them requires specifying how prices are set, which is the first assumption this

[18]The competition-shifting instruments common in this literature are not used at this level of aggregation, since the number of potential entrants in a market changes only when a carrier adds or drops service at an endpoint city; Appendix B.1 reports that specification and shows that the estimate of ε is not sensitive to the difference.

level adds to the framework of Section 3.

Firm behavior and equilibrium. In each period and market, airlines maximize operating profits by setting prices for the set of routes they operate ($\mathcal{K}_{imt}$), taking as given demand conditions, costs, and their route networks. The pre-tax price vector chosen by airline i maximizes $\Pi_{imt} = \sum_{k \in \mathcal{K}_{imt}} (\tilde{p}_k - c_k - \tau e_k)\, s_k$, where $s_k = q_k / N_m$ is the market share of product k implied by consumer demand. The first-order condition for each product k is $s_k + \sum_{j \in \mathcal{K}_{imt}} (\tilde{p}_j - c_j - \tau e_j) \frac{\partial s_j}{\partial \tilde{p}_k} = 0$.

Stacking the first-order conditions of all firms, equilibrium pre-tax prices and shares, $\tilde{\boldsymbol{P}}_{\boldsymbol{mt}}$ and $\boldsymbol{S}_{\boldsymbol{mt}}$, satisfy a Nash-Bertrand equilibrium characterized by the system

$$\boldsymbol{\mu}_{\boldsymbol{mt}} \equiv \tilde{\boldsymbol{P}}_{\boldsymbol{mt}} - \boldsymbol{C}_{\boldsymbol{mt}} = -J_{mt}^{-1} \boldsymbol{S}_{\boldsymbol{mt}}, \tag{9}$$

where $\boldsymbol{\mu}_{\boldsymbol{mt}}$ is the vector of operating markups defined in Section 3.1, $\boldsymbol{C}_{\boldsymbol{mt}}$ is the vector of marginal costs inclusive of carbon tax payments, and J_{mt} is the Jacobian of shares with respect to prices, multiplied element-wise by an indicator matrix of common ownership.[19] Equation (9) defines markups conditional on shares and the Jacobian; but both depend on prices, so equilibrium prices themselves are defined only implicitly and are obtained by numerically solving a nonlinear system. As such, markups are equilibrium objects: own- and cross-price demand responses, cost differences across products, and the allocation of products to firms jointly determine $\boldsymbol{\mu}_{\boldsymbol{mt}}$, and a tax change moves them.

Demand responses. Together with a demand system, equation (9) implicitly defines equilibrium prices and quantities as functions of the tax. Totally differentiating the stacked first-order conditions $\Omega\left(\tilde{\boldsymbol{P}}_{\boldsymbol{mt}}, \tau\right) = \boldsymbol{0}$ together with the demand system yields

$$\frac{d\tilde{\boldsymbol{P}}_{\boldsymbol{mt}}}{d\tau} = -\left[\frac{\partial \Omega}{\partial \tilde{\boldsymbol{P}}_{\boldsymbol{mt}}}\right]^{-1} \frac{\partial \Omega}{\partial \tau}, \qquad \frac{dq_k}{d\tau} = N_m \sum_{j \in \mathcal{K}_{mt}} \frac{\partial s_k}{\partial \tilde{p}_j} \frac{d\tilde{p}_j}{d\tau}, \tag{10}$$

where $\mathcal{K}_{mt}$ is the set of all products in a market. These expressions define the objects that Levels 1 and 2 replace with an aggregation restriction to work around the challenge of fully capturing own- and cross-product effects. This setting allows for within-market reallocation, so quantities and emissions need not move in proportion at some aggregate level.

[19] Cell ij of the ownership matrix equals 1 if products i and j are offered by the same airline and 0 otherwise.

This characterization has an additional advantage that extends the analysis beyond marginal changes. Evaluating the optimal tax and the revenue-neutral substitution requires pass-through, elasticities, and markups away from the observed equilibrium, because both involve large tax changes. Weyl and Fabinger (2013) show that non-marginal welfare variation depends on the path of pass-through between equilibria. Specifying such a path with an ad hoc functional form, whether linear or isoelastic, is itself a structural assumption (Kleven, 2021), though one without a clear economic rationale.

A specified equilibrium model disciplines the path; in practice, counterfactuals are computed by re-solving the system (9) at each tax level rather than by integrating derivatives. For aviation, validated models are available: I adopt Nash-Bertrand competition with discrete-choice demand, a standard framework in the empirical literature on this sector (e.g., Berry et al., 2006; Berry & Jia, 2010; Aguirregabiria & Ho, 2012; Pagoni & Psaraki-Kalouptsidi, 2016). Importantly, the Nash-Bertrand assumption maintained treats pricing as static and non-cooperative. Tacit collusion in the airline industry would imply markups above the Nash-Bertrand benchmark, and evidence of coordinated behavior has been documented (Ciliberto et al., 2019). To the extent that observed prices embed a degree of coordination, the equilibrium markups recovered here might understate true markups, which would widen the markup wedge and reinforce the paper's central finding that the second-best tax falls below marginal damages.[20]

Demand specification. Demand follows a nested logit discrete choice model (McFadden, 1978). Consumer n's utility from product k is $u_{nk} = V_k + \nu_n(\lambda) + \lambda\epsilon_{nk}$, with mean utility $V_k \equiv X_k^D \beta^D - \alpha_m p_k + \xi_k$, where X_k^D is a vector of observed characteristics, p_k is the ticket price, ξ_k captures characteristics unobserved in the data, and the error structure is standard, with parameter $\lambda \in [0,1]$ governing the correlation of tastes within nests. All flights are grouped in a single nest, and the outside option—not flying, with mean utility normalized to zero—is the single choice in a separate nest.[21] Parameter $\alpha_m \equiv \alpha_0 + \alpha_1\left(Y_m - \bar{Y}\right)$ captures market-specific sensitivity to prices motivated by income differences, where Y_m is the per-capita income at the origin of market m and $\bar{Y}$ is the

[20] Appendix C.3 discusses alternative conduct assumptions and performs a bounding exercise by re-estimating the model with full collusive behavior.

[21] As discussed in Section 7, the carbon tax examined here induces rather homogeneous cost changes within a market, because emissions per available seat vary substantially more across markets than within them: over 85% of the variance in e_k lies between markets. For this reason, the critical substitution is among nests and this parsimonious demand model captures the main mechanism transparently. Appendix B.3 considers two alternative nesting specifications used in the literature and shows that they result in coefficients that violate economic theory.

mean per-capita income among the origins included in the sample.

Applying the standard inversion of Berry (1994) to the nested logit share equations yields the linear estimating equation

$$\ln s_k - \ln s_0 = X_k^D \beta^D - \alpha_0 p_k - \alpha_1 \left(Y_m - \bar{Y}\right) p_k + (1-\lambda) \ln s_{k|g} + \xi_k, \tag{11}$$

from which the parameters β^D, α_0, α_1, and λ are estimated. The vector X_k^D includes the following observed product characteristics: (i) service frequency in departures per week, (ii) number of stops, (iii) market distance and its square, (iv) travel distance added due to connections and its square, (v) percent of delayed departures in the previous quarter, (vi) number of destinations offered by the airline from the origin city, and (vii) fixed effects for airline, origin-by-quarter, and destination-by-quarter. Consumer surplus is calculated using the standard log-sum expression for nested logit (Train, 2009).

Supply specification. Using the estimated parameters $\widehat{\alpha}_m$ and $\widehat{\lambda}$, observed prices, and observed market shares, predicted marginal operating costs follow from manipulating the equilibrium system (9): $\widehat{\boldsymbol{C}}_{\boldsymbol{mt}} = \tilde{\boldsymbol{P}}_{\boldsymbol{mt}} + \widehat{J}_{mt}^{-1} \boldsymbol{S}_{\boldsymbol{mt}}$. These predicted costs are then used to estimate the supply side of the model. The specification of the cost equation builds on the cost structure of flight operations, which maps into five categories according to their unit of variation: costs per block hour (aircraft operating costs proportional to usage time, plus passenger service costs), costs per departure (aircraft servicing and ground operations), costs per enplaned passenger (traffic servicing, such as passenger and baggage processing), costs per distance (primarily fuel), and indirect and overhead costs (Belobaba et al., 2015). The cost equation is

$$\widehat{c}_k = \rho F_k + \beta_i^S \text{Ramp-to-ramp}_k + \gamma_{i,o} + \gamma_{i,c_1} + \gamma_{i,d} + \gamma_{i,c_2} + \gamma_t + \omega_k, \tag{12}$$

where F_k is fuel expenditure per available seat; Ramp-to-ramp$_k$ is the flight duration measured in hours; $\gamma_{i,o}, \gamma_{i,c_1}, \gamma_{i,d}$, and γ_{i,c_2} are fixed effects for each airport along a route interacted with airline; γ_t is a quarter fixed effect; and ω_k is an idiosyncratic cost shock. The key parameter is ρ, which maps varying levels of fuel expenditure to implied costs. Ramp-to-ramp$_k$ proxies costs per block hour, with β_i^S accommodating differences across airlines; route-based fixed effects capture how average costs per enplaned passenger and per departure vary for each airline at each airport; and the time fixed effect captures average variation in costs across quarters.

Output units and load factors. Two margins of adjustment underlie the quantity of air travel supplied. Airlines choose how many seats to offer in discrete increments of aircraft and departures, while consumers choose how many trips to purchase. The ratio between seats occupied by paying consumers and available seats offered is the load factor. The model adopted here holds the load factor fixed: quantities count revenue passengers, marginal costs and emission intensities are measured per available seat, and the two units therefore remain proportional. This approximation allows for a continuous quantity to represent an output supplied in discrete increments, and it is consistent with how carriers manage capacity, since a single flight serves demand from every market connecting through its endpoints and carriers allocate seats across those markets to sustain a target occupancy. Appendix D documents the corresponding stability of observed load factors despite large shifts in fuel costs comparable to the ones simulated in this paper. This approximation, however, imposes a limitation: the model cannot capture the composition of an output reduction between fewer flights and emptier ones, and in the non-marginal scenarios of Section 7.2 it overlooks the discrete costs of adjusting fleet, crew, and ground operations that a substantial resizing of the network would involve.

Estimation. The joint demand-supply system is nonlinear in α_m and λ, so the rich set of fixed effects—corresponding to approximately 3,000 dummy variables—cannot be directly factored out of the equations, which raises memory requirements and computation time for GMM estimation. I address this limitation using the method proposed by Conlon and Gortmaker (2020), which absorbs nonlinear terms into the left-hand side, factors out fixed effects by alternating projections (Bauschke et al., 2003), and expresses linear parameters as functions of the nonlinear ones, so that the numerical optimization searches only over a two-dimensional space.

Identification. Three variables in equation (11) are potentially correlated with unobserved characteristics ξ_k: prices, within-nest shares, and flight frequency (Berry & Jia, 2010). To address this endogeneity, I construct instruments following the aviation literature (Berry et al., 2006; Berry & Jia, 2010; Aguirregabiria & Ho, 2012; Pagoni & Psaraki-Kalouptsidi, 2016). There are four groups of instruments. First, the competition-shifting instruments include (i) the number of airlines in a market, (ii) the number of products offered by competitors, (iii) the share of competitors' products that are nonstop flights, (iv) the number of potential legacy entrants, and (v) the number of potential low-cost entrants. Second, (vi) fuel cost per available seat is a cost-shifter. Third, (vii) complementary density along segments measures the number of passengers from other markets transported on the same segments of a route; this instrument indicates the scale of operations that are complementary to a product and affect both costs and frequency of service. The fourth group includes all exogenous variables in equation (11).

The exogeneity of fuel expenditures in (12) relies on three factors. First, jet fuel closely follows crude oil price shifts, as shown in Figure A-7. Hence, since jet fuel accounts for a small fraction of refined oil products, shocks to jet fuel prices largely reflect changes that are exogenous to the aviation sector.[22] Second, fleet and network composition affecting fuel use reflect long-term decisions that are unlikely to respond quickly to recent fuel price shocks; concerns over responses to small shocks are mitigated by the fact that the data set considers only a short period. Third, while fuel-saving operations are technically possible, these have narrow margins because they typically extend travel time and increase time-related operating costs, such as crew assignment and aircraft turnover (Belobaba et al., 2015).[23]

Given the above, identification relies on the exogeneity of the instruments in each equation, which can be expressed in vector moment conditions $E\left(Z_k^D \xi_k\right) = 0$ and $E\left(Z_k^S \omega_k\right) = 0$, where Z_k^D is the vector of demand instruments and Z_k^S is the vector of supply instruments (the vector of supply covariates, in this case). After factoring out fixed effects, these conditions form a system with 22 moment conditions and 19 parameters, the base of the two-step GMM estimation used in this paper.[24]

5.4 Comparing the levels of analysis

Table 2 places the three implementations side by side. Reading down the columns shows what each additional layer of structure costs, and reading across the rows shows what each contributes. The

[22] Another threat to identification could be fuel hedging strategies, which may shift effective fuel costs and bias estimates. In 2018, however, this channel had a very limited impact. Annual 10-K filings for the six largest carriers show that American and United did not hedge fuel costs, while Southwest held substantial hedging positions, and the other carriers held mixed and more limited hedging instruments. Nevertheless, hedging led to a maximum change in effective fuel costs of only 1.2% (net gains for Delta when combining returns on financial instruments and the Monroe refinery ownership) in comparison with moderate jet fuel spot price shifts (13% difference between the lowest and the highest quarterly average prices).

[23] For instance, Gosnell, List, and Metcalfe (2020) find statistically significant efficiency gains from monitoring and incentivizing pilots directly, but the potential effects show reductions of less than 1% in total fuel use for international flights; shorter distances in domestic flights offer less room for route adjustments, so those gains could be even smaller here. Appendix C.1 analyzes an alternative setting with input substitution and finds evidence that jet fuel has a very low elasticity of substitution.

[24] Appendix B.2 reports the results for an alternative strategy that estimates demand and supply separately using two-stage least squares.

first three rows describe the restrictions, the next two the inputs, and the last one the scope of questions each level can address.

The comparison makes two points that organize the results. First, the wedges and the SGMAC require far less than the welfare effects in levels: at Levels 1 and 2 they follow from accounting averages and observed market weights, with no estimated parameter entering. Second, the boundary that matters most is between Levels 2 and 3. Moving from the sector to the market relaxes an aggregation restriction and delivers cross-market heterogeneity at the cost of one estimated elasticity. Moving from the market to the product changes the character of the exercise: markups and demand responses become equilibrium objects that adjust with the tax, which allows for within-market reallocation and, more importantly, evaluation away from the observed equilibrium. Section 7 illustrates how coarser levels might lead to significant differences when analyzing large tax changes.

Table 2: Three levels of aggregation: assumptions, requirements, and scope.

	Level 1: Sector	**Level 2: Market**	**Level 3: Product**
Product definition	Single composite product	One composite per market	Airline-route offerings
Markups	Sector accounting average (Form 41, per ASM)	Firm-level accounting averages (Form 41, per ASM), scaled by distance	Equilibrium objects implied by Nash-Bertrand pricing
Demand response $\frac{dq_k}{d\tau}$	Cancels out; not required	Proportional within market; aggregate elasticity ε across markets	Implicit differentiation of the equilibrium system; counterfactual re-solution
Estimated parameters	None	ε (2SLS)	Full demand and supply system (GMM)
Key assumptions	Representative product; accounting markups approximate true markups	Complete marginal pass-through; within-market shares invariant to marginal changes; accounting markups approximate true markups	Nested logit demand; Nash-Bertrand conduct; constant marginal costs
Inputs for wedges and SGMAC	Four accounting averages	Accounting averages and observed market weights (no estimation)	Estimated equilibrium objects
Inputs for welfare effects in levels	Not available	Adds estimated ε	Estimated equilibrium objects
Delivers	Aggregate wedges and SGMAC	Adds cross-market heterogeneity; marginal welfare effects	Adds within-market reallocation; non-marginal reforms; optimal tax; tax substitution

6 Estimation results and model validation

Two of the three levels developed in Section 5 require estimated parameters: the market level requires the elasticity of aggregate demand, and the product level requires the full demand and supply system. This section reports the estimates for these levels and assesses how well the structural model reproduces patterns it was not estimated on.

6.1 The elasticity of aggregate demand

Table 3 shows the results of estimating equation (8) with Ordinary Least Squares (OLS) and Two-Stage Least Squares (2SLS). While results from both estimators indicate negative elasticities, 2SLS estimates are larger in absolute value and in line with the literature. This estimate is similar to Berry and Jia (2010), one of the few papers reporting aggregate price elasticities. Based on the estimated structural parameters of an industry model using 2006 data, they calculate a sector-aggregate elasticity of -1.67. As Section 5.2 shows, this elasticity does not enter the market-level wedges or the SGMAC, which depend only on observed market weights. However, it is required to express marginal changes in welfare, quantities, and emissions in levels.

Table 3: Estimation of aggregate price elasticity.

	OLS ln(Passengers)	2SLS ln(Passengers)
ln(Price) [ε]	-0.472	-1.848
	(0.055)	(0.138)
Fixed effects		
City pair	Yes	Yes
Quarter	Yes	Yes
Observations	19,750	19,750
R^2	0.963	0.959

Notes: Standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). Full regression results are reported in Table A-1 .

6.2 Structural parameters

Table 4 reports the results for the joint estimation procedure for key parameters. The main demand coefficients are largely in line with the literature. Here, $\alpha_0 = 0.936$ captures price sensitivity for a location with sample mean personal income per capita of about \$59,000 in 2018.[25] Parameter α_1 maps how price sensitivity varies with regional income. Combined, these two parameters allow price sensitivity to vary across markets, ranging from 0.53 in the highest-income origin (San Francisco, CA) to 1.17 in the lowest (El Paso, TX). Estimates for α in more recent studies generally vary between 1.36 (Aguirregabiria & Ho, 2012) and 0.45 (Pagoni & Psaraki-Kalouptsidi, 2016). Berry and Jia (2010) differentiate leisure and business travelers; based on 2006 data, they estimate the price parameter as 1.05 for the first group (estimated as 63% of consumers) and 0.10 for the second group. Parameter λ depends on the nesting assumption; among papers that also group all flights in the same nest, estimates of $1 - \lambda$ are typically between 0.3 and 0.4.

The supply side parameter shows that a \$1 variation observed in fuel costs per available seat translates, on average, to a \$0.745 change in implied costs. It is worth mentioning that ρ is not itself a cost pass-through parameter but, instead, a parameter that flexibly captures how marginal costs used for pricing vary with observed fuel costs. In equilibrium, the realized cost pass-through also depends on each market's demand and structure. Other parameters, reported in Table A-2, capture additional drivers of heterogeneity in consumer preferences and firm costs and also align with the literature. For instance, consumers value a higher frequency of service, meaning more opportunities for convenient travel times, and dislike connecting flights; legacy carriers have generally higher costs per block hour and passenger than low-cost ones.

6.3 Model validation

To assess the validity of the model, I compare predicted outcomes with out-of-sample values reported by airlines in the Form 41; these are the same reported averages that serve as the inputs to the sector-level calculation in Section 5.1. These data are not used to estimate structural parameters, so they provide a sanity check on the model. In particular, I evaluate the model's ability to generate reported patterns in average revenues, costs, and markups per ASM by airline, which are metrics are typically used in the aviation sector to analyze airline performance. For this paper, in particular, they capture a key point of the model's application: markups.

Table 5 shows how model outcomes compare with reported financials. The last row shows that predictions for revenue, cost, and markups reasonably approximate the values reported by airlines.

[25] In comparison, BEA national aggregates report a personal income per capita of \$53,111 for that year. Not surprisingly, this shows that the sample of locations with mid-size or larger airports skews towards areas of higher income.

Table 4: Estimates of key structural parameters.

Demand	
Price (\$100) $[-\boldsymbol{\alpha_0}]$	-0.936
	(0.116)
Price × Demeaned income $[-\boldsymbol{\alpha_1}]$	0.100
	(0.054)
ln(share within nest) $[\mathbf{1}-\boldsymbol{\lambda}]$	0.443
	(0.049)
Supply	
Fuel expenses $[\boldsymbol{\rho}]$	0.745
	(0.003)
Observations	267,967
Objective function minimum	1.33×10^{-5}
Hansen J-stat	3.56 (p = 0.31)

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The demand equation includes fixed effects for airline, origin airport-by-quarter, and destination airport-by-quarter. The supply equation includes fixed effects for quarter and for each route airport-by-airline. Estimates for the complete set of structural parameters are reported in Table A-2.

For individual airlines, however, the quality of the predictions varies. Differences in predicted versus reported revenues have two main explanations. The first reason is related to selection: the data set used for prediction skews towards larger and more competitive markets and, thus, may not accurately represent the average revenue in the whole network. The second reason is that the model does not capture product unbundling, so predicted revenue comes from ticket sales only. In practice, baggage, reservation, and cancellation fees make up a small fraction of airline revenues; however, for low-cost carriers (LCCs), these sources of revenue can represent a large share of total operating revenues (Belobaba et al., 2015; Brueckner et al., 2015). As a consequence, the model has limited ability to reproduce some business practices of LCCs and is less accurate in predicting their revenues. Note, also, that model predictions capture economic costs rather than accounting-based ones, such as the operating costs reported in Form 41; significant opportunity costs could further widen the gap between these two cost metrics.

Errors in revenue prediction, however, are not fully passed onto markup prediction differences because the model predicts costs that rationalize pricing choices via the Nash-Bertrand equilibria. Especially for LCCs, predicted average costs are substantially lower than reported figures. Nevertheless, these predictions result in markups that are closer to the reported values. Hence, even

Table 5: Average operating revenue, costs, and markups per available seat-mile (ASM).

	Revenue (cents/ASM)		Cost (cents/ASM)		Markup (cents/ASM)		Market share (%)
Airline	Predicted	Reported	Predicted	Reported	Predicted	Reported	
Southwest	13.49	13.70	9.99	9.44	3.49	4.26	27.88
American	14.57	14.65	11.83	10.37	2.74	4.28	19.26
Delta	15.91	15.80	12.97	10.89	2.94	4.92	18.93
United	13.91	12.60	11.33	10.75	2.58	1.85	14.30
Other LCCs	4.21	8.31	1.77	6.69	2.44	1.62	8.39
JetBlue	10.73	12.07	8.06	9.46	2.67	2.62	5.62
Alaska	10.07	11.48	7.15	8.37	2.93	3.11	5.61
Sector average	13.03	13.04	10.12	9.58	2.91	3.46	

Notes: Predicted averages result from the estimated structural model. Reported values are calculated based on the Form 41 Financial database. Market shares are based on total passengers enplaned in the period. Averages statistics are calculated based on sector aggregate revenues, costs, and ASM.

though these equilibria may not capture all components of pricing decisions, they reproduce the most important patterns in the reported data and provide a good approximation of sector averages.

The comparison also has a second use. The reported averages in Table 5 are the inputs to the sector-level calculation of Section 5.1, and the predicted averages are the aggregation of the product-level equilibrium markups. The two columns therefore bound how much the accounting approximation to markups differs from the equilibrium one at the sector aggregate.

7 Second-best carbon taxation

Based on the estimated sufficient statistics and structural parameters, this section performs counterfactual analyses to quantify the potential welfare effects of carbon taxation in the US domestic aviation sector. First, section 7.1 examines the local efficiency of carbon taxation by estimating marginal effects at the three levels of aggregation and tracing how the differences between them originate in the measurement of markups. Next, section 7.2 focuses on non-marginal changes to calculate optimal taxation levels under market distortions. Lastly, section 7.3 considers the effects of substituting the current sales tax for a revenue-neutral carbon tax.

The welfare analyses below consider three values for the SCC to estimate damages of carbon emissions. The low SCC scenario is set at \$50/tCO_2, a reference value commonly used in the literature and in line with the US Federal SCC of \$51 set in 2021. The medium SCC of \$200/tCO_2 rounds up recent estimates of \$185 (Rennert et al., 2022) and \$190 (EPA, 2022) based on a 2% discount rate. The high SCC scenario is set at \$300/tCO_2, rounding the estimate of \$308 in Rennert et al. (2022) when a 1.5% discount rate is used.

Two features of the joint distribution of fares and emission intensities organize the results in this section. First, emission intensity is primarily a market attribute: more than four-fifths of the variance in e_k lies between markets (85%, or 92% when products are weighted by passengers), so a carbon tax changes the relative cost of flying in different markets far more than it reorders the products within a market. Second, the cost shocks induced by the taxes considered here are moderate relative to fares. The average ticket embeds 0.42 tCO_2 (weighting products by passengers), so a \$100/tCO_2 tax adds \$42 to its cost, which corresponds to 11% of the passenger-weighted mean fare of \$385. In contrast, the within-market standard deviation of this burden corresponds to only about 6% of the within-market standard deviation of fares. As a consequence, the aggregate responses estimated below are driven primarily by travelers leaving the most affected markets (the substitution toward the outside option) rather than by reallocation among products within a market when this margin is available.

7.1 Social marginal abatement costs and welfare consequences

Table 6 shows the estimated effects of a marginal increase in the carbon tax at the three levels of aggregation developed in Section 5. The first column reports the sector-level calculation, whose wedges are ratios of accounting aggregates and require no estimation. This level is silent about quantity responses: projecting them requires an elasticity of aggregate demand, which sector aggregates (four quarterly observations) cannot credibly provide. The second column presents the market-level sufficient statistics, and the third the outcomes of solving the estimated structural model with a \$1/tCO_2 tax increase; for reference, this is equivalent to an additional jet fuel tax of 1.34 cents per gallon.

Despite the substantial differences in the data and assumptions underlying each level, the three implementations deliver SGMAC estimates of comparable magnitude.[26] The sector-level calculation yields \$228.53/tCO_2, which falls between the market-level estimate of \$242.51 and the structural estimate of \$196.36—a value 19% below the market-level one. The decomposition of the SGMACs shows that all three levels produce similar sales tax wedges, so the differences concentrate almost

[26] The SGMACs and welfare effects are calculated assuming a fixed fuel intensity. I relax this assumption in Appendix C.1, extending the model and using an estimated elasticity of substitution between fuel and non-fuel inputs to obtain bounds for the SGMAC. The results show that, for a plausible range of elasticities, adding input substitution results in modest reductions (8–13%) in the SGMAC estimates shown in Table 6.

entirely in the markup wedge.[27]

Table 6: Comparison of marginal effect estimates across methods.

Method	Accounting averages	Sufficient statistics	Structural parameters
Aggregation level	*Sector*	*Market*	*Product*
Marginal costs (\$/tCO$_2$)			
Markup wedge	175.61	185.12	139.34
Sales tax wedge	49.64	54.10	53.74
Baseline (equivalent) carbon tax	3.28	3.28	3.28
Social gross marginal abatement cost (SGMAC)	228.53	242.51	196.36
Marginal changes on market aggregates (per tCO$_2$)			
Mean ticket price	—	+0.12%	+0.09%
Passengers	—	−0.22%	−0.30%
Emissions	—	−0.24%	−0.37%

Differences in markup-wedges can be mapped to how each level measures markups, and the pattern is visible directly in the data. Figure 1 plots the mean markup of each decile of emission intensity under the two measurement approaches, weighting products by passengers. The accounting-based markups of the market-level implementation rise almost linearly with intensity, from \$24 in the least carbon-intensive decile to \$163 in the most intensive one. This gradient is a consequence of the measurement rather than of pricing behavior: accounting markups per seat-mile are available only at the carrier level and allocating each carrier's margin in proportion to distance makes markups nearly proportional to fuel use. As a result, this aggregation leads to a correlation between markups and emission intensity of 0.93.[28] The structural markups, recovered from observed pricing behavior, display no such association: their decile means remain within \$71–83 across the entire intensity distribution, and their correlation with emission intensity is close to zero. The contrast separates a specification assumption from an observed feature of the data. This difference is consequential for welfare consequences because the markup wedge weights markups by each product's contribution to the marginal emission response, so markups that rise with intensity

[27]The direction of this gap is consistent with conduct beyond Nash-Bertrand: the sector-average markup reported in financial data lies between the Nash-Bertrand and full-coordination predictions of the structural model (see Appendix C.3). Re-estimating the model under full coordination raises the SGMAC to \$260/tCO$_2$, with the increase due entirely to the markup wedge, while the qualitative conclusions are unchanged.

[28]This correlation is passenger-weighted; the unweighted correlation is 0.91. Using a uniform markup per seat-mile set at the sector-level mean instead of carrier-specific rates produces an essentially identical profile, with a correlation of 0.97. For reference, the correlation between the structural markups and emission intensity is −0.07 passenger-weighted and −0.17 unweighted.

receive the largest weights.

Figure 1: Markups and carbon intensity under alternative measurement approaches.

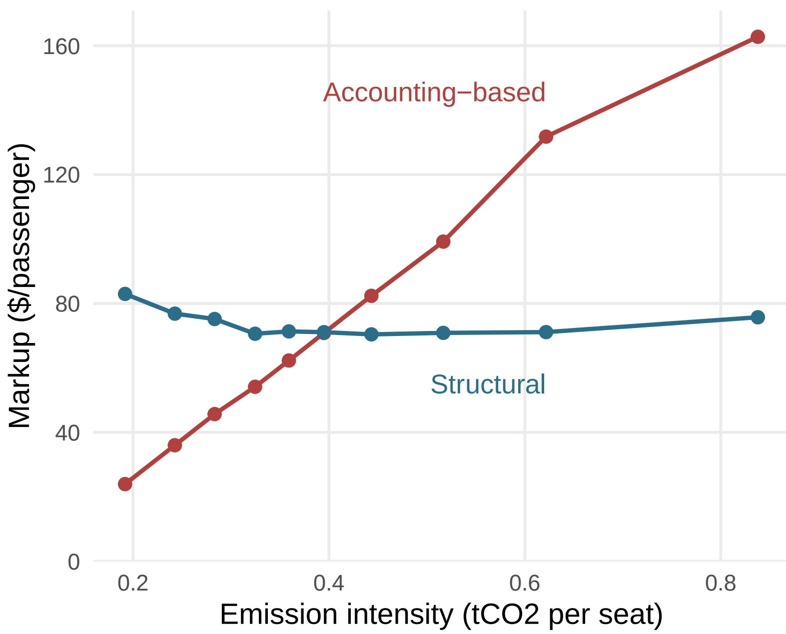

Notes: Each point is a decile of the distribution of emission intensity e_k across markets, weighted by baseline traffic, plotted at each decile mean intensity. *Accounting-based* markups allocate each carrier's Form 41 margin per available seat-mile in proportion to distance, as in the market-level implementation of Section 5.2. *Structural* markups are the passenger-weighted market means of the equilibrium markups estimated in Section 5.3.

The $45.78 gap between the market-level and structural markup wedges can be roughly decomposed by re-evaluating the market-level wedge using alternative markup imputations. First, replacing carrier-specific margin rates with the single sector-wide rate yields a wedge of $180.94; as such, carrier-level accounting variation explains little of the gap, and moves the wedge away from the structural benchmark. Second, rescaling the accounting markups to match the mean level of the structural ones lowers the wedge only from $185.12 to $177.20, because the two measurements imply similar average markups ($76.91 against $73.62 per passenger). Third, plugging in structurally estimated markups, aggregated to the market-level, reduces the wedge to $158.32. The remaining distance to the structural wedge of $139.34 reflects the aggregation restrictions—a common demand elasticity and no substitution within markets. Hence, of the total difference, roughly $8 is attributable to the level of accounting markups, $19 to their correlation with emission intensity, and $19 to aggregation. The accounting approximation discussed in Section 5.2 therefore biases the SGMAC upward mainly by forcing markups to rise with carbon intensity, and the direction of the bias is unambiguous: with markups measured this way, the market-level implementation overstates the marginal cost of abatement by about a fifth.

Market structure is a second dimension along which markups could vary systematically, and its role can be gauged from the relationship between structural markups and the number of competitors. Table 7 reports baseline markups by haul group and number of carriers.

Two patterns stand out. First, there is a monopoly premium, especially in short-haul markets, where market distance is essentially constant across competition bins: mean markups per mile[29] fall from 43 cents/mi under monopoly to 31.5 with three carriers and change little thereafter. Medium-haul markets show the same compression from 17.7 cents/mi in markets with up to two carriers to less than 10 cents beyond 4 players. Second, past this point the relationship is flat: among medium-haul markets between 1,000 and 2,300 miles, where distance is comparable across bins, per-mile markups are essentially constant from two carriers onward. In the spirit of Bresnahan and Reiss (1991), most of the association between concentration and margins is exhausted by the third or fourth competitor. The number of carriers, however, is not a sufficient statistic for markups in this setting. Carriers differ in business models and cost structures (differences visible in both their financial reports and the estimated supply parameters) and in the quality consumers attach to their products, so markup levels vary across markets with the same carrier count. For this reason, the welfare analysis relies on the estimated product-level markups rather than on market concentration summaries, and competition enters the results through the markups themselves, as equation (3) prescribes.

7.2 Non-marginal welfare changes and the optimal carbon tax

The marginal estimates of Table 6 leave open how outcomes scale under cost shocks of realistic magnitude. Between 2015 and 2018, monthly average jet fuel spot prices ranged from their minimum (January 2016) to their maximum (October 2018) by \$1.32 per gallon, which is approximately equivalent to \$98/tCO$_2$. Motivated by this variation, I evaluate the projected effects of raising the carbon tax from its baseline equivalent of \$3.28 to \$100/tCO$_2$—an increment comparable to the cost variation the sector has absorbed. This exercise also delimits the range over which each level of analysis remains informative. The sector level stops at the margin itself, since extrapolation requires an elasticity it cannot estimate. The market level extrapolates under its maintained assumptions: fuel cost increases are fully passed through to prices market by market, quantities respond through

[29] Markups are normalized by the distance between market endpoints for comparability. Since a direct flight is the shortest distance served, this denominator is common to all products in a market and does not vary with routing.

Table 7: Mean baseline markups per haul category and number of competitors.

Haul	Carriers	Passengers (M)	Share of pax (%)	City pairs	Markup (¢/pax-mi)	Mean market distance (mi)
Short	1	0.2	0.1	56	43.0	238
	2	0.9	0.6	57	42.1	221
	3	1.7	1.3	44	31.5	238
	4	1.9	1.4	31	31.5	243
	5+	1.8	1.3	8	30.3	244
Medium	1–2	1.4	1.0	277	17.7	565
	3	3.8	2.8	604	14.2	635
	4	19.6	14.4	1,076	12.5	720
	5	28.1	20.6	782	9.7	897
	6	31.1	22.8	435	8.4	1,058
	7	29.6	21.7	168	8.9	1,131
	8+	8.0	5.9	47	7.5	1,502
Long	1–4	0.3	0.2	57	2.9	2,455
	5–7	1.8	1.3	87	3.1	2,478
	8+	6.3	4.6	30	3.2	2,482

Notes: Mean baseline Nash-Bertrand markups are calculated using the estimated structural parameters and weighted by passengers. Haul groups are based on the distance between market endpoints and follow the US Environmental Protection Agency suggested classification: short haul $\leq$ 300 miles, medium haul $>$ 300 and $\leq$ 2,300 miles (EPA, 2018). Per-mile markups are normalized by market distance. City pairs counts are nondirectional. Cells pool carrier counts where markets are sparse.

the isoelastic form with the common elasticity ε, and markups, and hence the wedges, remain fixed at their baseline values. Only the structural model produces an equilibrium response. Table 8 compares the two extrapolations.

The comparison quantifies the cost of evaluating a large tax increase with statistics estimated at the baseline. The market-level extrapolation predicts a price increase of 11.3% against 9.1% in the structural model, since it imposes full pass-through while equilibrium markups partially absorb the tax. The bias reverses for quantities: passengers fall 17.7% under the extrapolation, compared with 24.7% in the structural model, because holding the elasticity at its baseline value understates how demand responds as prices rise and burdens concentrate in the most affected markets. Emissions fall faster than passengers at both levels, as composition effects across markets operate even in the extrapolation, but the gap is substantial: 19.2% against 29.0%, or 11.0 against 16.7 million tons abated. The resulting social gross average abatement cost, at \$428/tCO_2, overstates the structural

Table 8: Comparison of welfare effeccts of a \$100/tCO_2 carbon tax with market-level extrapolation and structural model.

	Sufficient statistics	Structural model
Mean ticket price	+11.3%	+9.1%
Passengers	−17.7%	−24.7%
Emissions	−19.2%	−29.0%
Emissions abated (Mt CO_2)	11.0	16.7
Social gross average abatement cost (\$/tCO_2)	427.9	238.0

Notes: The tax rises from its baseline equivalent of \$3.28 to \$100/tCO_2. The sufficient-statistics column extrapolates under the maintained market-level assumptions: full pass-through of fuel costs market by market, isoelastic quantity responses with the common elasticity ε, and markups fixed at baseline. Average abatement cost is the change in private surplus divided by the change in emissions. The sector level is not reported, as extrapolation requires an elasticity that sector aggregates cannot provide.

estimate of \$238 by nearly 80%. As such, marginal statistics remain informative near the baseline, but the optimal-tax and tax-substitution analyses that follow rely exclusively on the structural model.

The structural model makes it possible to evaluate non-marginal increases in the jet fuel tax and to recalculate the equilibrium outcomes in each market. All counterfactuals hold the set of competitors in each market fixed.[30] First, under a low SCC, social welfare decreases monotonically over a wide range of tax values, so no positive optimal tax exists when the SCC is \$50. Next, in the mid SCC scenario, the optimal tax is about \$7/tCO_2. As the current jet fuel tax is equivalent to \$3.28/tCO_2, this finding indicates that the welfare-maximizing carbon tax would be equivalent to a modest increase in the current jet fuel tax, raising it from 4.4 cents/gallon to about 9 cents/gallon; for reference, the mean jet fuel spot price in the period examined was \$2.02/gallon. Lastly, in a high SCC scenario, the optimal tax is approximately \$114/tCO_2, corresponding to a jet fuel tax of \$1.53/gallon. In both the mid and high SCC cases, the optimal taxes are much lower than the standard Pigouvian tax prescription of \$200 and \$300. The welfare curves underlying these optima, together with the corresponding marginal damage and abatement cost schedules, are reported in Appendix B.4.

To trace how market distortions produce this result, Figure 2 decomposes the SGMAC across

[30] A non-marginal carbon tax could in principle alter market structure by inducing carrier exit, a margin held fixed in this short-run analysis. Appendix C.2 gauges the impact of this channel. Reduced-form estimates indicate that carrier presence responds to fuel costs, but the implied effect is far smaller than the loss of a full market competitor. As a conservative upper bound, I simulate the optimal tax while forcing the exit of an entire competitor from every non-monopoly market. This exercise finds that the average abatement cost rises by roughly a third relative to the no-exit estimates.

taxation efforts. Per equation (2), the SGMAC is the sum of the carbon tax itself and the two distortion wedges, so the vertical composition of the stacked areas shows the sources of the gap between the second-best tax and marginal damages as the tax rises. At the baseline, the markup wedge accounts for about 72% of the combined distortions. As the tax increases, incomplete pass-through compresses equilibrium markups, shrinking the markup wedge, while more expensive tickets raise the revenue collected by the *ad valorem* sales tax and widen its wedge. At the high-SCC optimal tax of \$114, the wedge between the \$300 marginal damage and the optimal tax is about \$186, of which roughly \$124 is due to market power and \$62 to the sales tax. Market power is therefore the dominant distortion at every taxation effort considered, although its share declines as the tax accounts for a larger part of the restriction in output. The upper boundary of the stack also shows that the SGMAC itself rises steadily with the tax level, the marginal counterpart of the increasing average abatement cost documented in Table 8.

Figure 2: The composition of the social gross marginal abatement cost across tax levels.

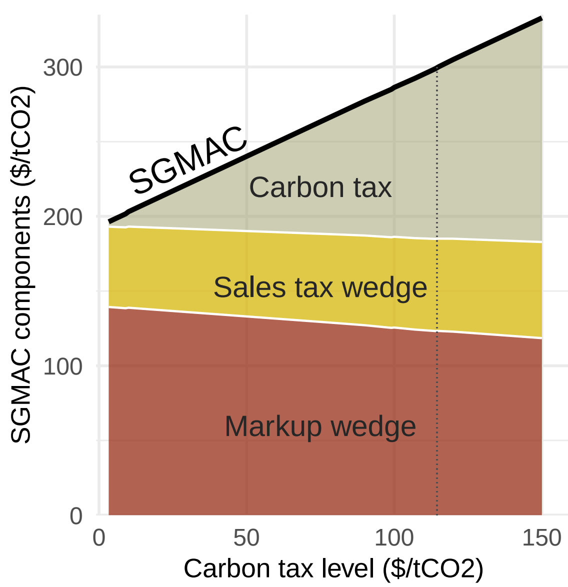

Notes: Each area is a component of the social gross marginal abatement cost (SGMAC) at the corresponding tax level, computed from the structural model's equilibrium at that tax: the markup wedge (bottom), the sales tax wedge (middle), and the carbon tax itself (top). The vertical dotted line indicates the high-SCC optimal tax level at \$114. The upper boundary of the stack is the SGMAC. The baseline tax is equivalent to \$3.28/tCO_2.

The private cost summarized by the SGMAC can be further decomposed into its incidence. Along the entire path of tax increases, the burden splits into roughly constant shares, with consumers bearing slightly more than half of the private losses. Under the high-SCC optimal tax of \$114/tCO_2 tax, consumer surplus falls by \$4.2 billion and operating profits by \$3.8 billion, with tax revenues increasing by \$3.5 billion that could be allocated to partially offset losses on either side. The cumulative incidence path is reported in Appendix B.4.

7.3 Is the double dividend possible?

Even though market distortions act in the direction of reducing aggregate emissions, they are imperfect substitutes for an externality tax, since sales taxes do not necessarily match the externalities generated by the carbon emissions of each product. For example, all else equal, consumers are willing to pay a premium for shorter nonstop flights, thus leading to higher equilibrium fares and sales taxes paid. If this premium exceeds the cost differentials to stop flights, then the nonstop flight would pay higher taxes per emission unit despite being more fuel-efficient and less polluting.

Welfare theory suggests that an alternative design could improve the efficiency of taxation as a climate policy instrument. A feasible alternative would still need to raise taxes, as the revenue from the existing sales tax funds several government operations, including the Federal Aviation Administration (see section 2). In this sense, replacing the sales tax with a revenue-neutral carbon tax that raises the same amount of revenue could pay a double dividend. First, it could decrease carbon emissions. Second, it could promote net welfare gains by replacing a less efficient taxation scheme. This substitution would, in theory, exactly offset any fiscal externalities and maintain the level of public funds raised (Hendren & Sprung-Keyser, 2020). However, as established in section 3.4, we cannot determine a priori whether both dividends materialize, as they depend on the particular distribution of prices and emissions in a sector.

To find the revenue-neutral carbon tax level, I run counterfactual simulations setting the sales tax to zero and incrementing the jet fuel tax by tenths of a cent. The resulting revenue-neutral fuel tax of \$0.829/gallon corresponds to a carbon tax of \$61.87/tCO_2.

Table 9 summarizes the changes in the counterfactual scenario implementing the revenue-neutral carbon tax. The top part of the table shows that such a tax substitution would result in virtually no abatement. Despite the mean ticket price staying relatively unaffected, a revenue-neutral carbon tax would increase passenger traffic by almost 1%, while emissions would fall by only 0.08%. The bottom part of the table confirms that a revenue-neutral carbon tax increases welfare in all three scenarios. However, these welfare improvements are driven by higher private surplus, most of which is due to a consumer surplus increase of \$113.9 M. The remainder accounts for gains in producer surplus, which occur despite the minimal decrease in mean markups. As such, these gains emerge primarily from eliminating the sales tax distortion; further calculations show that the markup wedge at the revenue-neutral tax is \$137.67, almost the same as in the baseline case.

At first glance, the results in Table 9 run counter to the intuition that a tax instrument proportional to the externality should substantially reduce emissions. Two mechanisms jointly produce this outcome: the burdens reallocated by the substitution are concentrated across markets rather than within them, and equilibrium markup adjustments absorb most of the reallocation. Figure 3 illustrates these mechanisms with the effects on flights within each quintile of carbon intensity, measured in CO_2 emissions per passenger.

Table 9: Effects of substituting the existing sales tax with a revenue-neutral carbon tax.

Aggregate market changes		
Mean ticket price		+0.47%
Passengers [million]		+0.94% [+1.29]
Emissions [million tons]		−0.08% [−0.05]
Mean markup [\$]		−1.12% [−0.50]
Welfare changes (million \$)		
Short-run private surplus (SRPS)		+176.6
Social welfare, SCC of:	\$50	+178.9
	\$200	+185.8
	\$300	+190.4

Panel (a) in Figure 3 shows the distribution of carbon taxes implied by the existing sales tax, which was defined as $r\tilde{p}_k/e_k$ in section 3.4.[31] These distributions display the potential welfare gains from tax substitution: sales taxes overtax emissions in the least carbon-intense flights while undertaxing emissions in the most carbon-intense options.[32] Since carbon intensity is primarily a market attribute, this mispricing operates across markets, so the burden reallocation that corrects it shifts costs from short-haul toward long-haul travel. A revenue-neutral carbon tax enforces a uniform tax of \$61.87/tCO_2 on all emissions, inducing ticket prices to adjust accordingly.

Unlike marginal changes outlined in the theoretical framework, new equilibrium prices here also respond to non-marginal effective tax changes on competing products. As such, market power allows firms to adjust markups in order to preserve or gain market share in response to changes in the tax

[31] Appendix B.5 presents additional information on the heterogeneity across carbon intensity quintiles. It shows that higher quintiles on average have longer distances between endpoints and higher prices, but no consistent differences in markups or estimated quality.

[32] This mispricing is the environmental counterpart of the redistribution that Miravete et al. (2020) document for uniform commodity taxation, where a single rate applied to differentiated products overprices some and underprices others. The relevant dimension here is carbon intensity rather than demand elasticity, and the implicit transfers are to the most carbon-intensive flights.

Figure 3: The effects of a revenue-neutral carbon tax substitution by carbon intensity quintiles.

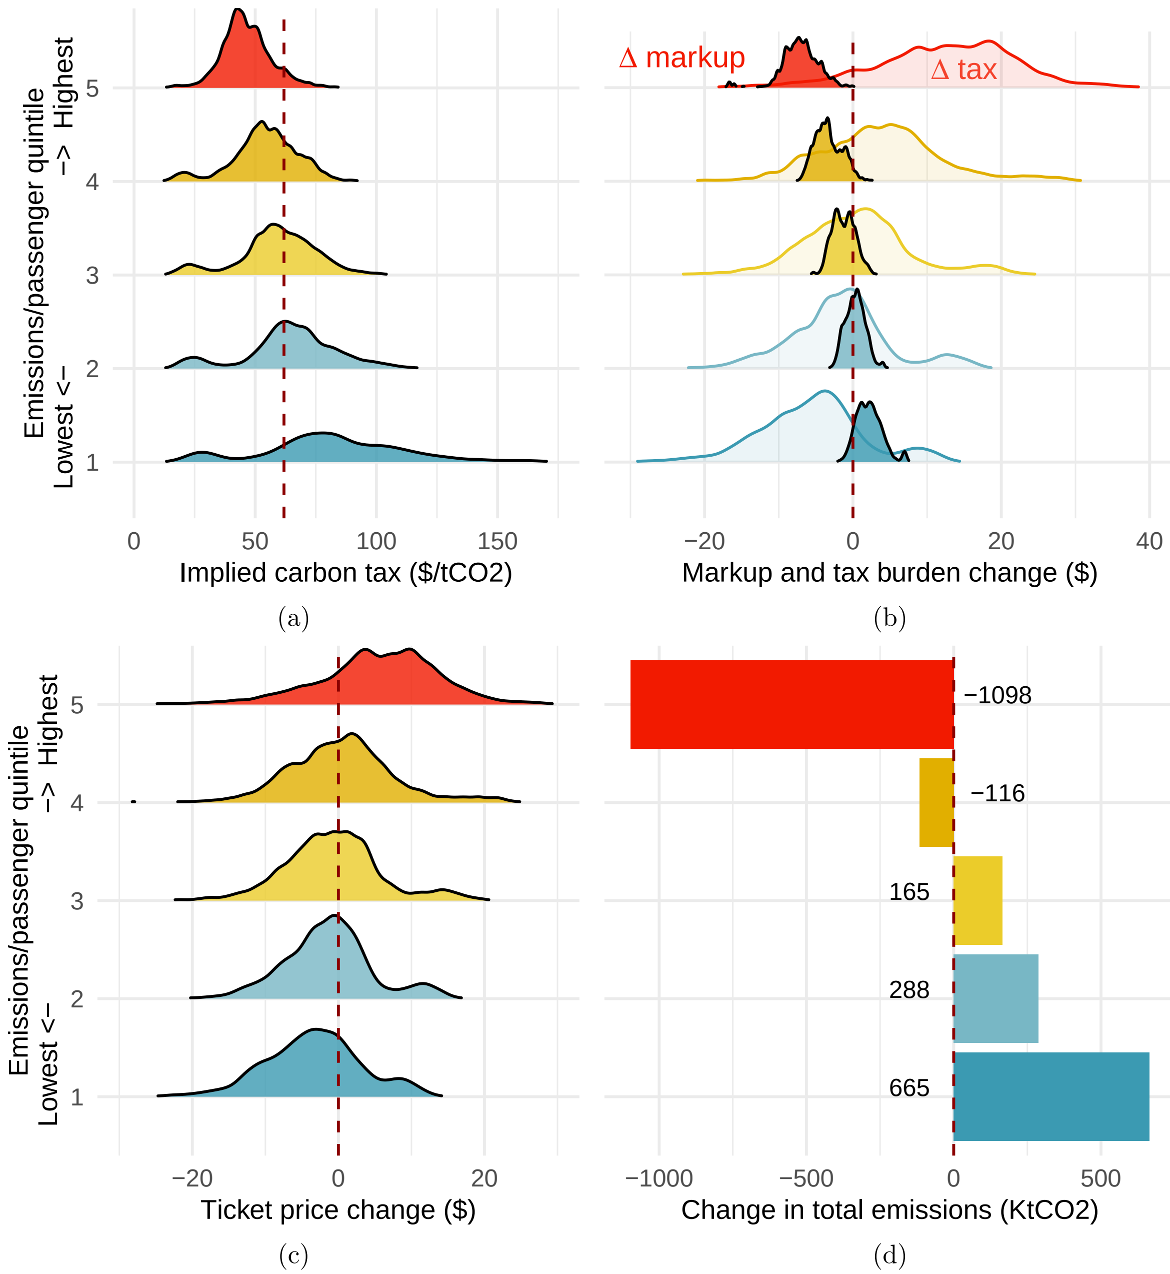

Notes: Carbon intensity is measured in CO_2 emissions per passenger in a given flight. Curves show kernel densities for each quintile of carbon intensity. Distributions are weighted by total emissions, so each quintile accounts for approximately the same aggregate emissions in the baseline. Quintile cut-off levels are approximately 320, 410, 530, and 710 Kg of CO_2 per passenger.

burden of each product. Panel (b) in Figure 3 shows the direction of these adjustments: markups fall for the vast majority of flights in the highest carbon-intensity quintile, with the opposite happening in the lowest quintile. The adjustments are small on average but systematically ordered by intensity. Weighting products by baseline traffic, the mean markup change is $-\$0.50$ per passenger; weighting by baseline emissions, it is −\$0.85. The difference between the two means indicates that markup reductions concentrate precisely where emissions concentrate, dampening the price increases that the substitution would otherwise produce for carbon-intensive products.

To isolate the role of these adjustments, I recompute the substitution counterfactual holding markups fixed at their baseline values, so that tax changes pass through mechanically and quantities respond through demand alone. This results in a slight increase in the revenue-neutral tax level, at \$63.28/tCO_2.[33] In this counterfactual, the substitution reduces aggregate emissions by 2.28% (1.31 million tons): the burden reallocation alone would deliver a modest environmental dividend, as demand falls in the carbon-intensive markets whose taxes rise. Restoring the equilibrium response eliminates almost all of this reduction, leaving the 0.08% decrease of Table 9. Hence, markup adjustments absorb roughly 96% of the environmental dividend that the reallocation of tax burdens would otherwise produce. The dividend at stake is itself small, as the mispricing corrected by the substitution is bounded by the dispersion of implied carbon taxes in panel (a), but the equilibrium response reduces it to approximately zero. Panel (c) in Figure 3 shows the consequence for prices: due to markup adjustments, prices fall for a substantial share of emissions in the upper quintiles. The combination of lower markups and the elimination of the deadweight loss of the sales tax contributes to increases in private surplus, as reported in Table 9. However, markup reductions in the upper quintiles shift the price change distribution to the left. As a result, the sharp decrease in emissions in the 5th quintile and the modest one in the 4th barely offset aggregate increases in the other quintiles, as shown in Panel (d).

This exercise shows that only one of the dividends materializes, as aggregate emissions remain virtually unchanged. Yet, a revenue-neutral tax substitution leads to non-negligible welfare gains as a result of private surplus increases following the removal of the distortive sales tax.

8 Conclusion

This paper has studied how oligopoly market power and pre-existing distortionary taxes affect the economic efficiency of carbon taxation. To do so, it incorporates the accounting of environmental damages into the welfare analysis of taxation under imperfect competition, and it connects this

[33] This difference relates to Miravete et al. (2018), who show that the tax rate and firms' pre-tax prices are strategic substitutes and trace the consequences for the revenue a given rate can raise. The instrument and the objective differ, as that paper contrast a regulator who anticipates the pricing response with one who does not, whereas the counterfactual here suppresses markup responses to isolate their contribution.

extended framework to empirics at three nested levels of aggregation with increasingly detailed data: a sector-level calculation based on reported accounting aggregates, market-level sufficient statistics, and a product-level structural model.

Despite the substantial differences in data requirements and assumptions, the three implementations deliver estimates of comparable magnitude, with social gross marginal abatement costs of \$229/tCO$_2$ at the sector level, \$243 at the market level, and \$196 in the structural model. Examining the differences between these implementations also quantifies the precision costs of the simplifying restrictions that less granular data impose. At the margin, approximating markups with accounting aggregates forces them to rise with carbon intensity, a pattern absent from the structural estimates, and this measurement alone leads the market-level implementation to overstate the marginal cost of abatement by about a fifth. This divergence increases when evaluating larger shocks, with the sufficient statistics approach reporting average abatement costs nearly 80% higher than those embedding the equilibrium responses of the structural model.

The main findings indicate that existing distortions in this sector are large: a marginal increase in the carbon tax reduces welfare whenever the social cost of carbon is below \$196/tCO$_2$. As a consequence, the optimal carbon tax is well below the standard Pigouvian prescription of pricing emissions at marginal damages. No positive optimal tax exists under a low SCC of \$50/tCO$_2$, while the optimum is \$6.76 under the medium value of \$200 and approximately \$114 under the high SCC value of \$300. Market power is the dominant distortion behind this gap, accounting for about 72% of the combined wedges at the baseline tax. Further analysis shows that substituting the existing sales tax with a revenue-neutral carbon tax would improve welfare but deliver virtually no abatement. This seemingly puzzling result stems from equilibrium markup adjustments: firms reduce markups on the most carbon-intensive flights, absorbing roughly 96% of the emission reductions that the reallocation of tax burdens would otherwise produce.

The analyses presented in this paper are subject to a number of limitations that also delineate directions for future research. First, all results concern short-run effects, as airlines' networks and fleets are held fixed. Higher fuel taxes may induce exit that increases the market power of remaining carriers, and the bounding exercise in Appendix C.2 indicates that this channel could raise abatement costs by up to a third. Future work could sharpen this bound with medium-run models in which market structure responds to taxation, as advances in the empirical analysis of dynamic network competition can make that margin tractable (Aguirregabiria & Ho, 2012). Second, all analyses are in partial equilibrium and overlook substitution toward other transportation modes, a leakage channel that the models used here do not capture. Embedding aviation within a demand system for multimodal travel, estimated with data that link trips across modes, would allow future studies to quantify this leakage and its welfare consequences. Third, load factors are held fixed, so the analysis treats a reduction in output as a proportional reduction in seats offered and passengers carried. Progress on this margin likely requires booking-level data connecting passenger itineraries to the occupancy of each flight, which would allow future work to relax the fixed load factor assumption and gauge the costs of discrete network downscaling. Fourth, climate change is the only externality considered, although fuel burn also affects local air pollution and health outcomes (Schlenker & Walker, 2016) and noise pollution lowers property values in the neighborhood of airports (Nelson, 2004). The welfare framework developed here accommodates additional externalities directly, so incorporating spatially resolved estimates of these local damages is a natural extension toward a more comprehensive accounting of the social costs of aviation.

With limitations considered, the results in this paper illustrate a crucial challenge for aviation: abatement via demand reduction has a high cost in terms of private welfare. Existing taxes and market power already drive equilibrium quantities down, so further reductions in demand come with significant private costs. Importantly, the estimated social abatement costs apply exclusively to the case of an aviation-only carbon tax without an economically-viable abatement technology. Alternative policies may, therefore, be more adequate in the medium run, including mandatory multi-sector emission permits or carbon offsets.[34] While emission reductions from an aviation-only carbon tax would only come from flying less, with permits and offsets firms can take advantage of lower abatement costs in other sectors to lower emissions. In practice, the emission charges from these alternative instruments would reduce net emissions both via demand reduction with higher ticket prices and abatement elsewhere, lowering the effective social gross marginal abatement cost of the policy. However, in the long run, more ambitious efforts to curb aviation emissions are likely to depend on technological advancements toward fuel alternatives, among which sustainable aviation fuels and electric and hydrogen-powered planes are potential candidates in development (Lohawala et al., 2026).

Lastly, the empirical exercise conducted here offers a methodological illustration that extends beyond aviation. The detail of publicly available data on this sector made it possible to implement the same welfare formulas at all three levels of aggregation, from a calculation requiring only

[34] Another option currently available is the offering of voluntary carbon offsets directly to customers. However, its promise to improve aviation sustainability has not materialized due to low participation rates and difficult verification of actual emission reductions (Mello, 2024).

reported accounting aggregates to a full structural model. In most settings, such granularity is unavailable, and welfare analyses of corrective taxation must rely on aggregate statistics and the simplifying assumptions they impose. Implementing the three levels side by side measures the approximation costs of those assumptions: coarse implementations remain informative for marginal effects, but their reliability deteriorates under non-marginal reforms and when the measurement of markups interacts with the distribution of emissions. These relative costs can serve as a guide for welfare analyses in sectors where only less detailed data are available.

References

Adachi, T., & Fabinger, M. (2022). Pass-through, welfare, and incidence under imperfect competition. *Journal of Public Economics*, *211*, 104589.

Aguirregabiria, V., & Ho, C.-Y. (2012). A dynamic oligopoly game of the US airline industry: Estimation and policy experiments. *Journal of Econometrics*, *168*(1), 156–173.

Anderson, S. P., de Palma, A., & Kreider, B. (2001). Tax incidence in differentiated product oligopoly. *Journal of Public Economics*, *81*(2), 173–192.

Aryal, G., Ciliberto, F., & Leyden, B. T. (2022). Coordinated capacity reductions and public communication in the airline industry. *The Review of Economic Studies*, *89*(6), 3055–3084.

Azar, C., & Johansson, D. J. (2012). Valuing the non-CO2 climate impacts of aviation. *Climatic Change*, *111*(3-4), 559–579.

Bauschke, H., Deutsch, F., Hundal, H., & Park, S.-H. (2003). Accelerating the convergence of the method of alternating projections. *Transactions of the American Mathematical Society*, *355*(9), 3433–3461.

BEA. (2018). *Regional Economic Accounts.* Bureau of Economic Analysis. Retrieved June, 2020, from `https://www.bea.gov/data/economic-accounts/regional`

Belobaba, P., Odoni, A., & Barnhart, C. (2015). *The global airline industry.* John Wiley & Sons.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.

Berry, S., Carnall, M., & Spiller, P. T. (2006). Airline hubs: costs, markups and the implications of customer heterogeneity. *Competition policy and antitrust*.

Berry, S., & Jia, P. (2010). Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*, *2*(3), 1–43.

Borenstein, S. (1989). Hubs and high fares: dominance and market power in the US airline industry. *The RAND Journal of Economics*, 344–365.

Borenstein, S. (1991). The dominant-firm advantage in multiproduct industries: Evidence from the us airlines. *The Quarterly Journal of Economics*, *106*(4), 1237–1266.

Borenstein, S., & Rose, N. L. (2014). How airline markets work... or do they? regulatory reform in the airline industry. In *Economic regulation and its reform: What have we learned?* (pp. 63–135). University of Chicago Press.

Bovenberg, A. L., & Goulder, L. H. (1996). Optimal environmental taxation in the presence of other taxes: General-equilibrium analyses. *American Economic Review*, *86*(4), 985-1000.

Bovenberg, A. L., & Mooij, R. A. D. (1994). Environmental levies and distortionary taxation. *The American Economic Review*, *84*(4), 1085–1089.

Bresnahan, T. F., & Reiss, P. C. (1991). Entry and competition in concentrated markets. *Journal of Political Economy*, *99*(5), 977–1009.

Brueckner, J. K., & Abreu, C. (2017). Airline fuel usage and carbon emissions: Determining factors. *Journal of Air Transport Management*, *62*, 10–17.

Brueckner, J. K., Lee, D., & Singer, E. S. (2013). Airline competition and domestic us airfares: A comprehensive reappraisal. *Economics of Transportation*, *2*(1), 1–17.

Brueckner, J. K., Lee, D. N., Picard, P. M., & Singer, E. (2015). Product unbundling in the travel industry: The economics of airline bag fees. *Journal of Economics & Management Strategy*, *24*(3), 457–484.

BTS. (2018). *Data products: Transtats.* Bureau of Transportation Statistics. Retrieved June, 2020, from `https://www.transtats.bts.gov`

Buchanan, J. M. (1969). External diseconomies, corrective taxes, and market structure. *American Economic Review*, *59*(1), 174–177.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, *1*(1), 451–488.

Ciliberto, F., & Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, *77*(6), 1791–1828.

Ciliberto, F., Watkins, E., & Williams, J. W. (2019). Collusive pricing patterns in the us airline industry. *International Journal of Industrial Organization*, *62*, 136–157.

Ciliberto, F., & Williams, J. W. (2010). Limited access to airport facilities and market power in the airline industry. *The Journal of Law and Economics*, *53*(3), 467–495.

Ciliberto, F., & Williams, J. W. (2014). Does multimarket contact facilitate tacit collusion? Inference on conduct parameters in the airline industry. *The RAND Journal of Economics*, *45*(4), 764–791.

Ciliberto, F., & Zhang, Z. (2017). Multiple equilibria and deterrence in airline markets. *Economic Inquiry*, *55*(1), 319–338.

Conlon, C., & Gortmaker, J. (2020). Best practices for differentiated products demand estimation with PyBLP. *The RAND Journal of Economics*, *51*(4), 1108–1161.

Cremer, H., Gahvari, F., & Ladoux, N. (1998). Externalities and optimal taxation. *Journal of Public Economics*, *70*(3), 343–364.

DOE. (2021). *Sustainable Aviation Fuel Grand Challenge.* Retrieved 2025-01-27, from `https://www.energy.gov/eere/bioenergy/sustainable-aviation-fuel-grand-challenge`

EIA. (2016). *Carbon dioxide emissions coefficients by fuel.* Energy Information Agency. Retrieved Dec. 12, 2019, from `https://www.eia.gov/environment/emissions/co2_vol_mass.php`

EIA. (2018). *Refiner Petroleum Product Prices by Sales Type.* U.S. Energy Information Administration. Retrieved June, 2020, from `https://www.eia.gov/dnav/pet/pet_pri_refoth_dcu_nus_m.htm`

EPA. (2018). *Greenhouse gas inventory guidance: Indirect emissions from events and conferences.* U.S. Environmental Protection Agency. Retrieved from `https://www.epa.gov/sites/default/files/2018-12/documents/indirectemissions_draft2_12212018_b_508pass_3.pdf`

EPA. (2022). *Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances.* US Environmental Protection Agency. Retrieved from `https://www.epa.gov/environmental-economics/scghg`

European Parliament. (2015). *Emission reduction targets for international aviation and shipping.* Directorate General for Internal Policies, European Parliament. Retrieved June, 2020, from `https://www.europarl.europa.eu/thinktank/en/document.html?reference=IPOL_STU(2015)569964`

Evans, W. N., & Kessides, I. N. (1994). Living by the "golden rule": Multimarket contact in the US airline industry. *The Quarterly Journal of Economics*, *109*(2), 341–366.

FAA. (2020). *Airport & Airway Trust Fund.* Federal Aviation Administration. Retrieved June, 2020, from `https://www.faa.gov/about/budget/aatf/`

FAA. (2021). *2021 United States Aviation Climate Action Plan.* Federal Aviation Administration. Retrieved August, 2023, from `https://www.faa.gov/sites/faa.gov/files/2021-11/Aviation_Climate_Action_Plan.pdf`

FCCC. (2018). *National greenhouse gas inventory data for the period 1990-2016.* Framework Convention on Climate Change, United Nations. Retrieved June, 2020, from `https://unfccc.int/sites/default/files/resource/17e.pdf`

Fukui, H., & Miyoshi, C. (2017). The impact of aviation fuel tax on fuel consumption and carbon emissions: The case of the us airline industry. *Transportation Research Part D: Transport and Environment*, *50*, 234–253.

Fullerton, D., & Metcalf, G. E. (2002). Cap and trade policies in the presence of monopoly and distortionary taxation. *Resource and energy economics*, *24*(4), 327–347.

Ganapati, S., Shapiro, J. S., & Walker, R. (2020). Energy cost pass-through in US manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, *12*(2), 303–342.

Genakos, C., Grey, F., & Ritz, R. (2020). *Generalized linear competition: From pass-through to policy* (Cambridge Working Paper in Economics No. 2078). Energy Policy Research Group, University of Cambridge. Retrieved from `http://www.jstor.org/stable/resrep30491`

Goolsbee, A., & Syverson, C. (2008). How do incumbents respond to the threat of entry? Evidence from the major airlines. *The Quarterly Journal of Economics*, *123*(4), 1611–1633.

Gosnell, G. K., List, J. A., & Metcalfe, R. D. (2020). The impact of management practices on employee productivity: A field experiment with airline captains. *Journal of Political Economy*, *128*(4), 1195–1233.

Goulder, L. H., Parry, I. W., Williams Iii, R. C., & Burtraw, D. (1999). The cost-effectiveness of alternative instruments for environmental protection in a second-best setting. *Journal of Public Economics*, *72*(3), 329–360.

Hendren, N., & Sprung-Keyser, B. (2020). A unified welfare analysis of government policies. *Quarterly Journal of Economics*, *135*(3), 1209–1318.

IATA. (2019). *World Air Transport Statistics.* International Air Transport Association. Retrieved June, 2020, from `https://www.iata.org/en/services/statistics`

ICAO. (2018). *Carbon Emissions Calculator Methodology, Version 11.* International Civil Aviation Organization. Retrieved June, 2020, from `https://www.icao.int/environmental-protection/CarbonOffset/Documents/Methodology%20ICAO%20Carbon%20Calculator_v11-2018.pdf`

Jacobsen, M. R., Knittel, C. R., Sallee, J. M., & Van Benthem, A. A. (2020). The use of regression statistics to analyze imperfect pricing policies. *Journal of Political Economy*, *128*(5), 1826–1876.

Kellogg, R., & Reguant, M. (2021). Energy and environmental markets, industrial organization, and regulation. In *Handbook of industrial organization* (Vol. 5, pp. 615–742). Elsevier.

Kim, E. H., & Singal, V. (1993). Mergers and market power: Evidence from the airline industry. *American Economic Review*, 549–569.

Kleven, H. J. (2021). Sufficient statistics revisited. *Annual Review of Economics*, *13*, 515–538.

Lade, G. E., & Bushnell, J. (2019). Fuel subsidy pass-through and market structure: Evidence from the renewable fuel standard. *Journal of the Association of Environmental and Resource Economists*, *6*(3), 563–592.

Lederman, M. (2007). Do enhancements to loyalty programs affect demand? The impact of international frequent flyer partnerships on domestic airline demand. *The RAND Journal of Economics*, *38*(4), 1134–1158.

Lederman, M. (2008). Are frequent-flyer programs a cause of the "hub premium"? *Journal of Economics & Management Strategy*, *17*(1), 35–66.

Lee, D., Fahey, D. W., Forster, P. M., Newton, P. J., Wit, R. C., Lim, L. L., … Sausen, R. (2009). Aviation and global climate change in the 21st century. *Atmospheric Environment*, *43*(22-23), 3520–3537.

Lee, D., & Luengo-Prado, M. J. (2005). The impact of passenger mix on reported" hub premiums" in the US airline industry. *Southern Economic Journal*, 372–394.

Lohawala, N., Toman, M., & Joiner, E. (2026). Sustainable aviation fuels for decarbonizing air travel. *Review of Environmental Economics and Policy*, *20*(2), 194–215.

Mansur, E. T. (2007). Upstream competition and vertical integration in electricity markets. *The Journal of Law and Economics*, *50*(1), 125–156.

McFadden, D. (1978). Modeling the choice of residential location. *Transportation Research Record*(673).

Mello, F. P. d. (2024). Voluntary carbon offset programs in aviation: A systematic literature review. *Transport Policy*, *147*, 158–168.

Miravete, E. J., Seim, K., & Thurk, J. (2018). Market power and the laffer curve. *Econometrica*, *86*(5), 1651–1687.

Miravete, E. J., Seim, K., & Thurk, J. (2020). One markup to rule them all: Taxation by liquor pricing regulation. *American Economic Journal: Microeconomics*, *12*(1), 1–41.

Morrison, S. A. (2001). Actual, adjacent, and potential competition estimating the full effect of Southwest Airlines. *Journal of Transport Economics and Policy*, *35*(2), 239–256.

Muehlegger, E., & Sweeney, R. L. (2022). Pass-through of own and rival cost shocks: Evidence from the us fracking boom. *Review of Economics and Statistics*, *104*(6), 1361–1369.

Nelson, J. P. (2004). Meta-analysis of airport noise and hedonic property values. *Journal of Transport Economics and Policy (JTEP)*, *38*(1), 1–27.

Owen, B., Lee, D. S., & Lim, L. (2010). Flying into the Future: Aviation Emissions Scenarios to 2050. *Environmental Science & Technology*, *44*(7), 2255-2260.

Pagoni, I., & Psaraki-Kalouptsidi, V. (2016). The impact of carbon emission fees on passenger demand and air fares: A game theoretic approach. *Journal of Air Transport Management*, *55*, 41–51.

Parry, I. W. (1995). Pollution taxes and revenue recycling. *Journal of Environmental Economics and Management*, *29*(3), S64–S77.

Phaneuf, D. J., & Requate, T. (2016). *A course in environmental economics: theory, policy, and practice.* Cambridge University Press.

Preonas, L. (2024). Market power in coal shipping and implications for us climate policy. *Review of Economic Studies*, *91*(4), 2508–2537.

Rennert, K., Errickson, F., Prest, B. C., Rennels, L., Newell, R. G., Pizer, W., … others (2022). Comprehensive evidence implies a higher social cost of co2. *Nature*, *610*(7933), 687–692.

Requate, T. (2006). Environmental policy under imperfect competition. *The International Yearbook of Environmental and Resource Economics*, *2007*, 120–207.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, *80*(3), 1019–1061.

Schlenker, W., & Walker, W. R. (2016). Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, *83*(2), 768–809.

Train, K. E. (2009). *Discrete choice methods with simulation.* Cambridge University Press.

Weyl, E. G., & Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, *121*(3), 528–583.

Winchester, N., Wollersheim, C., Clewlow, R., Jost, N. C., Paltsev, S., Reilly, J. M., & Waitz, I. A. (2013). The impact of climate policy on US aviation. *Journal of Transport Economics and Policy*, *47*(1), 1–15.

Appendix for

"Second-best carbon taxation in a differentiated oligopoly"

Diego S. Cardoso*

*University of Illinois Urbana-Champaign, Department of Agricultural and Consumer Economics. Email: dcardoso@illinois.edu.

Contents

A Welfare analysis derivations

This appendix derives the expressions underlying the welfare analysis in the main text. Appendix A.1 differentiates the four welfare components of Section 3 and obtains the marginal welfare expression (1). Appendix A.2 develops the revenue-neutral tax substitution results of Section 3.4. Appendix A.3 writes the market-level price and quantity responses in terms of the sufficient statistics of Section 5.2, which can then be substituted into the general formulas of Appendix A.1.

A.1 Marginal welfare derivations

Differentiating each welfare component defined in Section 3 shows how welfare in market m responds to a marginal change in the carbon tax τ. The tax is expressed in dollars per tCO_2 and implemented as a volumetric tax of τh per gallon of jet fuel, so tax payments per unit of product k equal τe_k. Since the fees ι_k are fixed, consumer prices satisfy $\frac{dp_k}{d\tau} = (1+r)\frac{d\tilde{p}_k}{d\tau}$. I assume that an equilibrium exists and is differentiable in a neighborhood of the current tax but remain agnostic about the specific price responses ($\frac{d\tilde{p}_k}{d\tau}$) and quantity responses ($\frac{dq_k}{d\tau}$). Differentiating each component with respect to τ (and omitting function arguments for clarity) yields

$$\frac{dCS_m}{d\tau} = -\sum_k q_k \frac{dp_k}{d\tau} \tag{A-1}$$

$$\frac{d\Pi_m}{d\tau} = \sum_k \left\{ \mu_k \frac{dq_k}{d\tau} + \left[\frac{d\tilde{p}_k}{d\tau} - e_k \right] q_k \right\} \tag{A-2}$$

$$\frac{dT_m}{d\tau} = \sum_k \left[q_k e_k + \tau e_k \frac{dq_k}{d\tau} \right] + r \sum_k \left[q_k \frac{d\tilde{p}_k}{d\tau} + \tilde{p}_k \frac{dq_k}{d\tau} \right] \tag{A-3}$$

$$\frac{d\Phi_m}{d\tau} = \phi \sum_k e_k \frac{dq_k}{d\tau}, \tag{A-4}$$

where sums run over the K_m products available in the market. Equation (A-1) uses the first-order optimality condition $u'_k(x^*_k) = \alpha_m p_k$ for every product k. In (A-2), $\mu_k \equiv \tilde{p}_k - c_k - \tau e_k$ denotes the operating markup, and $\frac{d\tilde{p}_k}{d\tau} - e_k$ captures its marginal change, which equals zero under complete pass-through of the carbon tax. In (A-3), the first sum represents carbon tax receipts, combining the mechanical payment $q_k e_k$ with the erosion of the tax base as quantities adjust. The second sum tracks sales tax receipts through the responses of both prices and quantities. In (A-4), ϕ is the marginal damage per tCO_2 and $e_k = hf_k$ converts fuel burn per available seat into emissions.

Summing the four components according to $W_m = CS_m + \Pi_m + T_m - \Phi_m$ and substituting

$\frac{dp_k}{d\tau} = (1+r)\,\frac{d\tilde{p}_k}{d\tau}$, two sets of terms cancel. First, the terms involving price derivatives: a marginal increase in the pre-tax price withdraws $(1+r)\,q_k \frac{d\tilde{p}_k}{d\tau}$ from consumer surplus and delivers exactly that amount to firms and the government combined, so these terms constitute a transfer among the three accounts. This cancellation uses only the consumer optimality condition; no firm first-order condition enters, and no restriction is placed on how markups adjust. Second, the mechanical carbon tax payment $q_k e_k$ enters operating profits negatively and tax revenue positively, so it is a transfer from firms to the government. The remaining terms involve only quantity responses,

$$\frac{dW_m}{d\tau} = \sum_k \left[r\tilde{p}_k + \mu_k + (\tau - \phi)\,e_k\right] \frac{dq_k}{d\tau}. \tag{A-5}$$

Since aggregate welfare is defined as $W \equiv \sum_m W_m$, summing (A-5) over markets recovers equation (1) in the main text, with the sum over k running over all products in the sector. The remaining derivations proceed as in Section 3.3: the ratio of the marginal change in private surplus to the marginal change in emissions yields the SGMAC in equation (2), and setting (1) to zero produces the second-best tax in equation (3).

A.2 Revenue-neutral tax substitution

This subsection derives the expressions in Section 3.4 and reports the restricted case that they nest. Sums run over all products in the sector, since both tax instruments are set nationally. Let $t_k \equiv \tau e_k + r\tilde{p}_k$, $E \equiv \sum_k q_k e_k$, and $B \equiv \sum_k q_k \tilde{p}_k$. Equilibrium quantities and pre-tax prices are functions of both instruments, $q_k(\tau, r)$ and $\tilde{p}_k(\tau, r)$, and I assume only that an equilibrium exists and is differentiable in a neighborhood of the current tax system.

Revenue neutrality. Differentiating $T = \sum_k q_k t_k$ with respect to τ and r gives equation (4), in which the first term is mechanical, the second collects the erosion of both tax bases as consumers adjust, and the third collects the change in the sales tax base induced by movements in pre-tax prices at constant quantities. This last term is the fiscal externality operating through markups: when firms absorb part of a tax change, the sales tax base moves even if no passenger switches. The quantity responses can be written in terms of the demand Jacobian and the pass-through of each instrument,

$$\frac{\partial q_k}{\partial \tau} = \sum_j \frac{\partial q_k}{\partial p_j} (1+r) \frac{\partial \tilde{p}_j}{\partial \tau}, \tag{A-6}$$

$$\frac{\partial q_k}{\partial r} = \sum_j \frac{\partial q_k}{\partial p_j} \left[(1+r) \frac{\partial \tilde{p}_j}{\partial r} + \tilde{p}_j \right], \tag{A-7}$$

which use $p_k = (1+r)\,\tilde{p}_k + \iota_k$. These expressions make explicit the two objects that a structural model supplies: the off-diagonal terms of $\frac{\partial q_k}{\partial p_j}$ and the pass-through vectors $\frac{\partial \tilde{p}}{\partial \tau}$ and $\frac{\partial \tilde{p}}{\partial r}$, both of which follow from differentiating the equilibrium system (9).

Marginal welfare along the revenue-neutral path. Totally differentiating each welfare component and using the consumer first-order condition $u_k'(x_k^*) = \alpha_m p_k$ yields

$$\begin{aligned}
dCS &= -\sum_k q_k dp_k, \\
d\Pi &= \sum_k \mu_k dq_k + \sum_k q_k d\tilde{p}_k - E\,d\tau, \\
dT &= \sum_k t_k dq_k + E\,d\tau + B\,dr + r\sum_k q_k d\tilde{p}_k, \\
d\Phi &= \phi \sum_k e_k dq_k,
\end{aligned}$$

with $\mu_k \equiv \tilde{p}_k - c_k - \tau e_k$. Substituting $dp_k = (1+r)\,d\tilde{p}_k + \tilde{p}_k dr$ into dCS and adding $d\Pi$, the price terms collapse to $-r\sum_k q_k d\tilde{p}_k - B\,dr$. In doing so, $dW = dCS + d\Pi + dT - d\Phi$ with dT set to zero becomes

$$dW = \sum_k (\mu_k - \phi e_k)\,dq_k - r\sum_k q_k d\tilde{p}_k - B\,dr - E\,d\tau.$$

The revenue constraint itself implies $-(E\,d\tau + B\,dr + r\sum_k q_k d\tilde{p}_k) = \sum_k t_k dq_k$. Substituting this identity delivers

$$\left.\frac{dW}{d\tau}\right|_{dT=0} = \sum_k \left[r\tilde{p}_k + \mu_k + (\tau - \phi)\,e_k\right] \frac{dq_k}{d\tau}, \tag{A-8}$$

which is equation (1) with $\frac{dq_k}{d\tau}$ evaluated along the revenue-neutral path. The derivation places no restriction on $\frac{\partial \tilde{p}_k}{\partial \tau}$ or $\frac{\partial \tilde{p}_k}{\partial r}$: the price-response terms enter the revenue constraint and the private-surplus components with opposite signs and cancel, exactly as in the fixed-r case. The revenue-neutral reform therefore alters the magnitude of the demand responses but leaves the set of objects requiring measurement unchanged.

Restricted case. To gain further intuition on the mechanisms, we can examine a simpler case. Suppose that (i) the emissions tax is completely passed through into pre-tax prices, $\frac{\partial \tilde{p}_k}{\partial \tau} = e_k$; (ii) pre-tax prices do not respond to the sales tax rate, $\frac{\partial \tilde{p}_k}{\partial r} = 0$; and (iii) demand for each product depends only on its own price, $\frac{\partial q_k}{\partial p_j} = 0$ for $j \neq k$. Conditions (i) and (ii) together state that markups are invariant to both instruments. Under these conditions (A-6)–(A-7) reduce to $\frac{\partial q_k}{\partial \tau} = (1+r)\, e_k \frac{dq_k}{dp_k}$ and $\frac{\partial q_k}{\partial r} = \tilde{p}_k \frac{dq_k}{dp_k}$, and (4) collapses to

$$\frac{dr}{d\tau} = -\,(1+r)\,\frac{\sum_k e_k \left[q_k + t_k \frac{dq_k}{dp_k}\right]}{\sum_k \tilde{p}_k \left[q_k + t_k \frac{dq_k}{dp_k}\right]}, \tag{A-9}$$

with associated demand response $\frac{dq_k}{d\tau} = \left[(1+r)\, e_k + \tilde{p}_k \frac{dr}{d\tau}\right] \frac{dq_k}{dp_k}$. Only in this case does the burden condition (3.4) translate directly into a sign for the quantity response: $\frac{dq_k}{d\tau} \leq 0$ if and only if θ_k is below the corresponding weighted average of implied taxes. Conditions (i) through (iii) correspond to the assumptions maintained by the sector- and market-level implementations in Section 5, but are not imposed by the structural implementation. Condition (ii) is restrictive under imperfect competition, since an *ad valorem* tax scales the margin rather than shifting it and is therefore partly absorbed by firms.

A.3 Market-level responses with sufficient statistics

This subsection expresses the price and quantity responses of the market composite, $\frac{dp_{mt}}{d\tau}$ and $\frac{dQ_{mt}}{d\tau}$, in terms of the sufficient statistics of Section 5.2. Consider the composite product $Q_{mt} = \sum_k q_k$ formed by the flights in market m at period t, where sums in this subsection run over the products in that market. The attributes of the composite are averages weighted by the within-market shares $s_{k|g}$: the pre-tax price $\tilde{p}_{mt} = \sum_k s_{k|g} \tilde{p}_k$, the price $p_{mt} = \sum_k s_{k|g} p_k$, the emission intensity $e_{mt} = \sum_k s_{k|g} e_k = h f_{mt}$, the fuel expenditure per available seat $F_{mt} = \sum_k s_{k|g} F_k$, and the markup $\mu_{mt} = \sum_k s_{k|g} \mu_k$. As discussed in Section 5.2, markups, non-fuel costs, and relative market shares are assumed invariant to a marginal change in the tax. Under these restrictions, a marginal increment in the carbon tax is fully passed through to pre-tax prices, $\frac{d\tilde{p}_k}{d\tau} = e_k$, and quantities move proportionally, $dq_k = s_{k|g}\, dQ_{mt}$.

The price response of the composite follows directly. Since $\frac{dp_k}{d\tau} = (1+r)\,\frac{d\tilde{p}_k}{d\tau}$ and shares are constant,

$$\frac{dp_{mt}}{d\tau} = \sum_k s_{k|g} \frac{dp_k}{d\tau} = (1+r)\, e_{mt} = \eta_{mt}\, p_{mt} \frac{e_{mt}}{F_{mt}}, \tag{A-10}$$

where the last equality uses the value of the pass-through elasticity under complete pass-through, $\eta_{mt} = (1+r)\, F_{mt}/p_{mt}$, and matches the product-level equation (6).

The quantity response of each product combines the full set of price responses, $\frac{dq_k}{d\tau} = \sum_j \frac{\partial q_k}{\partial p_j} \frac{dp_j}{d\tau}$. Summing over the products in the market and using $\frac{dp_j}{d\tau} = (1+r)\, e_j$ yields

$$\frac{dQ_{mt}}{d\tau} = (1+r) \sum_k \sum_j \frac{\partial q_k}{\partial p_j} e_j = (1+r) \sum_j e_j \frac{\partial Q_{mt}}{\partial p_j}. \tag{A-11}$$

With fixed relative shares, market-aggregate demand can be written as a function of the average price p_{mt} alone, and the weight of each product in that average is its share, $\frac{\partial p_{mt}}{\partial p_j} = s_{j|g}$. Hence,

$$\frac{\partial Q_{mt}}{\partial p_j} = \frac{\partial Q_{mt}}{\partial p_{mt}} \frac{\partial p_{mt}}{\partial p_j} = s_{j|g} \frac{\partial Q_{mt}}{\partial p_{mt}}. \tag{A-12}$$

Combining (A-11) and (A-12) and writing the result in terms of the elasticity of aggregate demand $\varepsilon = \frac{\partial \ln Q_{mt}}{\partial \ln p_{mt}}$ gives

$$\frac{dQ_{mt}}{d\tau} = (1+r)\, e_{mt} \frac{\partial Q_{mt}}{\partial p_{mt}} = (1+r)\, \frac{e_{mt}}{p_{mt}} Q_{mt} \varepsilon = \varepsilon\, \eta_{mt} \frac{e_{mt}}{F_{mt}} Q_{mt}, \tag{A-13}$$

which is equation (7) in the main text.

In the market-level aggregation, these two responses are the only unobserved inputs to the welfare formulas. Since quantities move proportionally, share-weighted sums collapse to composite terms, as in $\sum_k \mu_k \frac{dq_k}{d\tau} = \mu_{mt} \frac{dQ_{mt}}{d\tau}$, so market m enters the component derivatives of Appendix A.1 as a single product with attributes $(p_{mt}, \tilde{p}_{mt}, e_{mt}, \mu_{mt})$. Substituting (A-10) and (A-13) into (A-1)–(A-4) evaluates each welfare component at this level of aggregation; the wedges, the SGMAC, and the second-best tax then follow from the same steps as in the main text.

B Additional results and alternative estimations

The first three subsections of this appendix complement section 5 of the paper by showing complete regression tables, estimation results using an alternative estimator, and alternative demand specifications. The fourth subsection complements section 7.2, displaying additional panels and results representing the details of second-best carbon taxation levels. Finally, the fifth subsection complements section 7.3 by reporting additional results.

B.1 Full regression tables of main specifications

Table A-1 displays estimates for all parameters in the OLS and 2SLS specifications that estimate the market-aggregate price elasticity of demand, including the first-stage estimations. The middle pair of columns adds the numbers of potential legacy and potential low-cost entrants to the excluded instruments, following the aviation literature. These instruments are close to cross-sectional at this level of aggregation, since they change only when a carrier adds or drops service at an endpoint city, and they raise the within R^2 of the first stage from 0.115 only to 0.117 and lower the conditional F-statistic from 105.0 to 67.8, which raises concerns about weak instruments. The estimated elasticity moves from -1.848 to -1.797, well within a standard error, so the specification reported in the main text identifies ε from the cost shifter alone.

Table A-2 reports the full set of estimates of structural parameters. The additional parameters here reported display well-documented patterns, both on the demand and supply side. The estimates show that a larger number of offered destinations increases the value of frequent flier programs, thus making consumers more willing to travel with airlines that have more destination options. Moreover, the positive coefficient on market distance captures the value of air travel. The results also confirm that consumers dislike connecting flights, as they increase both the number of stops and the total distance traveled; the positive coefficient on extra distance squared indicates that this marginal disutility decreases with distance. Finally, consumers also dislike flights that were more frequently delayed in the previous quarter, although the effect is quite small. On the cost side, parameters also indicate an expected pattern: the costs per block hour and per passenger are generally greater among legacy carriers than in low-cost carriers.

B.2 Estimating the structural model with two-stage least squares

In this paper, the demand and supply sides of the structural model are jointly estimated using the generalized method of moments (GMM). An alternative approach used in Aguirregabiria and Ho (2012) is to estimate the demand and supply equations separately using two-stage least squares (2SLS). However, these two estimators are not equivalent. As Conlon and Gortmaker (2020) point out, estimating separate equations with 2SLS implicitly assumes that errors across equations are uncorrelated: $E[\xi_k \omega_k] = 0$. Even though this assumption may not hold in all applications, the 2SLS estimator is sometimes preferred because it is faster to estimate due to the linearity in parameters. In this section, I show the results of estimating the structural model using 2SLS and the first stage estimates when instrumental variables are used.

Table A-3 shows the estimates for the demand equation using 2SLS. This table shows that demand estimates using 2SLS and GMM (table 4) are generally similar but slightly smaller in absolute value. To provide further evidence of instrument validity, table A-4 shows the first-stage regressions for each instrumented variable, as well as the respective standard and conditional F-statistics.

Finally, table A-5 shows estimates of supply-side parameters. Here, estimates are slightly different from those from GMM estimation. In particular, the fuel cost parameter ρ is slightly smaller when estimated by 2SLS, though estimates are less precise in this case due to larger standard errors.

B.3 Alternative demand specifications

Section 5.3 shows that the predictions using the estimated structural model reproduce key patterns in out-of-sample data. Nevertheless, it is helpful to gauge whether alternative choices in the model would substantially affect the main parameters. Table A-6 considers plausible alternatives based on the literature.

As in Aguirregabiria and Ho (2012), specification (1) is estimated via OLS to illustrate endogeneity bias. When instruments are not used, the price coefficient becomes positive, and the nesting coefficient is negative—both violate standard assumptions of demand models. Specification (2) does not include high-dimensional fixed effects combining location and time. A comparison with the GMM estimates shows that α falls to roughly a quarter of its value, while other key parameters are less drastically affected. Specification (3) removes the interaction between prices and demeaned income, therefore resulting in a single price sensitivity parameter that indicates a 10% reduction in sensitivity relative to the main 2SLS regression.

Specifications (4) and (5) in Table A-6 consider alternative nesting choices. Recall that the main specification in the paper includes all flights in a single nest, whereas the outside option is kept in a separate nest. Specification (4) nests flights by airline, as in Aguirregabiria and Ho (2012). The results show that this specification does not seem to adjust well to the data, as its estimated price coefficient is positive, violating model assumptions. Since the distinction between stop and nonstop flights can be relevant for consumer choice, I also consider an alternative that nests these flight types separately in specification (5). However, as in the previous case, the results show that nesting by flight type violates model assumptions because λ is negative.

B.4 Additional results on the second-best carbon tax

This appendix reports the welfare schedules underlying the optimal taxes of Section 7.2. Here, *short-run private surplus* (SRPS) refers to the unweighted sum of the first three components. Panel (a) in Figure A-1 shows social welfare across tax levels for the three SCC scenarios: welfare declines monotonically under the low SCC and is single-peaked under the mid and high scenarios, so the optima reported in the main text are interior and unique. Panel (c) displays the same information in marginal terms, with the optimal tax at the intersection of marginal damages avoided and the marginal loss of short-run private surplus. Panel (b) plots the social gross marginal abatement cost and its wedge components against abatement levels, complementing Figure 2, which presents the same composition against tax levels. Panel (d) reports the cumulative incidence of tax increases on consumer surplus, operating profits, and tax revenue; the near-linearity of the three paths underlies the approximately constant burden shares noted in the main text.

B.5 Additional results on tax substitution

This appendix presents further information relative to the effects of replacing the existing sales tax with a revenue-neutral carbon tax. Figure A-2 displays the distribution of key characteristics of flights in the baseline (before any tax changes) in each of the carbon-intensity quintiles.

Panels (a) and (b) in Figure A-2 illustrate the crucial heterogeneity in products across quintiles: not surprisingly, higher emissions per passenger on average correspond to longer flights. Higher fuel costs for longer flights generally reflect higher pre-tax prices. However, a substantial share of the flights in each quintile can be found in the price range of \$250–500. Hence, these two panels provide further insights into the distribution of implied taxes depicted in panel (a) of Figure 3: the long right tail of implied taxes in the 1st and 2nd quintiles correspond to expensive but relatively short flights.

In contrast, panels (c) and (d) in Figure A-2 show that differences in baseline markups or quality are not relevant determinants in the heterogeneity of price responses. The distribution of markups varies little across the quintiles 1 through 4; the distribution of the 5th quintile, however, shows a slightly higher share of flights in the range of lower markups. Panel (d) shows the normalized distribution of $X_k^D \hat{\beta}^D$—the estimated mean utility of product characteristics, which can be interpreted as an index of product quality. We note that the dispersion in quality increases with the quintile, but no consistent differences in mean quality exist. The minor exception is again the 5th quintile: the mean index for quintiles 1 to 4 is between -0.7 and 0.7, but 0.35 for quintile 5. Nevertheless, these differences are relatively small and not sufficient to drive the results reported in Section 7.3.

C Bounding abatement costs with relaxed assumptions

C.1 Input substitution

The calculation of the effects of a marginal increase in the carbon tax hold constant the fuel burn per available seat. This assumption is a reasonable approximation for short-run impacts because there are limited options for fuel consumption adjustments, especially in short and medium-haul domestic flights. However, substantial fuel cost shocks can induce adjustments in aircraft configuration and cruise speed reductions—thus trading off fuel for travel time-based costs, such as labor.

This appendix extends the analysis of marginal effects in Section 7.1 to account for limited substitution of inputs as an exercise to provide lower bounds for the marginal abatement costs. To do so, I extend the model described in Appendix A to allow for input substitution by assuming that the production follows a Constant Elasticity of Substitution (CES) function. Then, I estimate the elasticity of substitution between fuel and non-fuel inputs using carrier costs reported in the Form P-5.2 data set. Using this elasticity, I re-calculate marginal effects and compare them with those shown in the main body of the paper.

C.1.1 Extending the model

This extension maintains the definition of a product as the combination of an airline, a route, and a quarter. In the main model used in this paper, the fuel amount used to supply an available seat for a given product is held constant. This assumption can be interpreted as a Leontief production function, where fuel cannot be substituted with other inputs. This extension will assume instead that the production function is a CES function, where fuel and non-fuel inputs can be substituted. Importantly, this derivation assumes that non-fuel inputs do not lead to any emissions.

The derivations in this appendix express the tax in volumetric terms, with τ denoting a tax per gallon of jet fuel and $\omega_k = w_k + \tau$ the tax-inclusive fuel price. This convention simplifies the notation because the substitution margin operates through the fuel price rather than through emission intensities. Since the carbon intensity of combustion h is a physical constant that input substitution does not alter, the correspondence of Section 3.1 applies unchanged. This is, a carbon tax of one dollar per tCO_2 corresponds to a volumetric tax of h dollars per gallon, marginal effects with respect to the carbon tax equal h times their volumetric counterparts, and the abatement cost expressions below are already stated per tCO_2 through the factor $1/h$.

The extended production function is given by

$$q_k = A_k \left(a_k f_k^{\frac{\sigma-1}{\sigma}} + (1-a_k)\, o_k^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{A-14}$$

where q_k is the quantity of available seats in product k, f_k is the fuel use in product k, o_k is a composite of non-fuel inputs, σ is the elasticity of substitution between fuel and non-fuel inputs, and A_k and a_k are CES parameters. The cost of o_k is normalized to 1 throughout, so the substitution captures the shift in cost shares. Moreover, it can be shown that share parameters can be calibrated so that they match the initial cost shares of fuel and non-fuel inputs: $a_k = \frac{\omega_k f_k}{\omega_k f_k + o_k}$.

From cost minimization, we obtain that the marginal unit cost is constant and given by

$$c_k = \frac{1}{A_k} \left(a_k^{\sigma} \omega_k^{1-\sigma} + (1-a_k)^{\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{A-15}$$

Moreover, a marginal change in the tax leads to a marginal change in unit costs proportional to fuel use: $\frac{dc_k}{d\tau} = f_k$. The marginal effects on prices and quantities, thus, remain identical to the baseline case. However, there will be an additional component for the marginal change in emissions due to a reduction in f; marginal effects on tax revenues, damages, welfare, and SGMAC will be different. It can be shown that the marginal change in fuel use following a marginal change in the tax is given by

$$\frac{df_k}{d\tau} = \frac{df_k}{d\omega_k} = -\sigma \frac{f_k}{\omega_k} (1-a_k). \tag{A-16}$$

Sufficient statistics approach. As discussed in Appendix A.3, this approach aggregates products in a given market to Q_{mt}. Given the fuel substitution characterized above, it can be shown that marginal changes to tax revenue are now given by

$$\frac{dT_{mt}}{d\tau} = Q_{mt}\left\{(1+r)\,f_{mt} + (\tau f_{mt} + r\tilde{p}_{mt})\,\frac{\eta_{mt}}{\omega_{mt}}\varepsilon\right\} - \tau\sigma\frac{Z_{mt}}{\omega_{mt}}Q_{mt}, \tag{A-17}$$

where $Z_{mt} = \sum_k s_{k|g} f_k\,(1-a_k)$ and the rightmost term now accounts for fuel tax revenue changes due to input substitution.

Similarly, marginal changes in damages are given by

$$\frac{d\Phi_{mt}}{d\tau} = \phi\left[\varepsilon\frac{\eta_{mt}}{\omega_{mt}} f_{mt} Q_{mt} - \sigma\frac{Z_{mt}}{\omega_{mt}} Q_m\right], \tag{A-18}$$

where the rightmost term captures reduced damages induced by lower fuel intensity.

The subsequent calculation of net welfare effects follows the same steps as in the baseline case. The resulting expression for the marginal abatement cost in a market at the baseline tax τ_0 is then given by

$$SGMAC_{mt}\,(\tau_0) = \frac{1}{h}\left[\tau_0 + \frac{\mu_{mt} + r\tilde{p}_{mt}}{f_m + \sigma\frac{Z_{mt}}{\eta_{mt}}}\right], \tag{A-19}$$

from where we note that increasing σ decreases the marginal abatement cost. However, the intensity of this reduction depends on the ratio $\frac{Z_{mt}}{\eta_{mt}}$.

Structural approach. The structural approach has two main differences relative to the baseline case. First, non-marginal cost adjustments will account for the lower fuel use per available seat. Second, since this extension captures fuel input substitution and its impact on marginal costs, there will be no parameter ρ representing how fuel expenditures map into costs. The derivation for this part uses exact algebra notation (Dekle, Eaton, & Kortum, 2008), where $\hat{x} \equiv x_{\text{post}}/x_{\text{pre}}$ represents the ratio of the post to pre-shock values of some variable x.

An increase in the jet fuel tax from τ_0 to τ_1 can be represented as $\hat{\omega}_k = (w_k + \tau_1)\,/\,(w_k + \tau_0)$. Then, it can be shown that the marginal cost adjust as follows

$$\hat{c}_k = \left(1 + a_k\left(\hat{\omega}_k{}^{1-\sigma} - 1\right)\right)^{\frac{1}{1-\sigma}}. \tag{A-20}$$

Moreover, the change is fuel use is given by

$$\hat{f}_k = \left(\frac{\hat{c}_k}{\hat{\omega}_k}\right)^{\sigma}. \tag{A-21}$$

Equations A-20 and A-21 are sufficient to simulate the equilibria with fuel use adjustment under different levels of the tax; the calculation of the marginal effects on welfare and the SGMAC follows as in the baseline.

C.1.2 Estimating the elasticity of fuel substitution

To estimate the elasticity of substitution, I use the data set of Form 41 Schedule T-2, which reports quarterly traffic and fuel use, among other variables, by aircraft model and airline. This allows me to calculate fuel use per available seat-mile for each aircraft model and airline. For this estimation, the period of analysis is extended to include all quarters from 2014 to 2018. Table A-7 shows summary statistics for the estimation data set.

The estimation of the elasticity of substitution is guided by equation (A-16), with three estimating equations considered: linear, log-log, and the CES-implied formula. The specifications are as follows:

$$f_k = \psi_{linear}\omega_k + \nu_p + \nu_i + \zeta_k, \tag{A-22}$$

$$\log\left(f_k\right) = \psi_{log}\log\left(\omega_k\right) + \nu_p + \nu_i + \zeta_k, \tag{A-23}$$

$$\log\left(f_k\right) = \sigma\log\left(\omega_k\right) \times \left(1 - a_k\right) + \nu_p + \nu_i + \zeta_k, \tag{A-24}$$

where f_k is fuel use per hundred available seat-mile, ω_k is the fuel cost in dollars per gallon, ν_p and ν_i are aircraft model and airline fixed effects, and ζ_k is an error term. The estimated coefficients of the linear (ψ_{linear}) or log-log (ψ_{log}) specifications can be used to calculate the implied mean elasticity of substitution using equation (A-16) and mean observed values of $\left(1 - a_k\right) f_k/\omega_k$ or $\left(1 - a_k\right)$. The coefficient estimated in the CES-implied formula directly maps to σ.

The identification of the ψ_{linear} and ψ_{log} relies on the exogeneity of jet fuel prices, which was discussed in the cost specification in Section 5.3. The identification of the CES-implied formula is more challenging because the share parameter a_k is a function of both fuel and non-fuel input prices. Nevertheless, the results are informative about the plausible range of the elasticity of substitution.

Table A-8 shows the estimates of the elasticity of substitution for the three specifications. The linear and log-log estimates imply a low elasticity of substitution, while the CES-implied estimate suggests a slightly higher value. Despite their differences, all three specifications point to the conclusion that the elasticity of substitution is relatively low. This result is consistent with the observation that there are limited alternatives to reduce fuel use in the short run (Belobaba, Odoni, & Barnhart, 2015).

C.1.3 Marginal abatement costs and welfare with input substitution

The expressions derived for the extended model and the empirical estimates of σ in the previous subsections are used to calculate the effects of a marginal increase in the carbon tax. The goal of this exercise is to establish reasonable bounds for these effects if fuel substitution were considered. In that spirit, the following quantification adopts two alternative levels for σ: 0.04 and 0.08, which are approximately the lower and upper bounds of the 95% confidence interval of the higher, CES-implied estimate.

Table A-9 shows the results of raising the jet fuel tax by 1 cent/gallon under different levels of fuel substitutability. For the sufficient statistics approach, increasing σ leads to the same aggregate effects in prices and traffic, but with reduced emissions. As a consequence, the SGMAC and the wedges are smaller. In the case of a higher substitutability ($\sigma = 0.08$), the SGMAC is about 13% lower than the baseline estimate; at $\approx$ \$212, the SGMAC gets closer to the value obtained for the baseline structural approach.

The structural approach yields similar results. It is important to note, however, that the case of $\sigma = 0$ here is *not* the same as the baseline because the extension version does not include the ρ parameter mapping fuel expenditures to implied marginal costs. Therefore, the effective pass-through in this case is higher than the baseline case without input substitution, and we observe more pronounced responses in prices and traffic. Nevertheless, the resulting SGMAC for the case of $\sigma = 0$ here is very similar to the baseline estimates. Increasing σ also leads to reductions in the SGMAC and wedges. In particular, the higher sigma results in a SGMAC that is about 8% lower that the baseline estimate or the $\sigma = 0$ cases.

These results above are consistent with the intuition that input substitution allows for reduced emission intensity, which translates into a lower abatement cost. However, using empirical estimates of the elasticity of substitution between fuel and non-fuel inputs, the magnitude of the reduction in SGMAC is relatively small. These bounding exercises indicate that the main results of the paper are a reasonable approximation of the effects of a carbon tax increase, even when input substitution is considered.

C.2 Tax-induced changes in market structure

The counterfactual analyses in section 7 hold the set of competitors in each market fixed. A non-marginal carbon tax, however, raises operating costs and could in principle induce some carriers to stop serving a market, thus altering market structure and the markups that drive the second-best results. This appendix assesses the robustness of the main findings to this channel in two steps. First, I provide reduced-form evidence that carrier presence responds to fuel costs, using historical variation in jet fuel prices. Second, I take the estimated structural model as a steady state representation and simulate the optimal tax while forcing the exit of a competitor in each market, therefore constructing a deliberately conservative upper bound on the additional cost of abatement.

This exercise does not attempt to estimate a structural model of entry and exit. Entry in this industry occurs at the level of the airport segment rather than the product or market, so an airline expands by adding service from an airport adjacent to its existing operations. Entry and exit are therefore the outcome of a dynamic game played over the entire route network, the state space of which is intractable for the US system unless significant simplifications on the strategic choice set and demand-supply models are made; estimating such a game requires strong simplifying assumptions, and the resulting counterfactuals are further complicated by equilibrium multiplicity (Aguirregabiria & Ho, 2012). Moreover, exit decisions in aviation reflect strategic and institutional considerations beyond short-run profitability, including entry deterrence and the occupancy of constrained gates and slots (Belobaba et al., 2015). The bounding approach below avoids these difficulties while still quantifying how much the main results could possibly change.

C.2.1 Reduced-form evidence

I construct a segment-level panel covering the sixteen quarters from 2015 through 2018. The unit of observation is a directional airport-pair (segment) in a given quarter. For each segment-quarter, I count the number of distinct carrier groups operating at least one departure, merging small and regional carriers into their controlling airline as in the main analysis. Operations data are drawn from Table T-100 of the Form 41 Traffic Database (BTS, 2018), and the national jet fuel spot price is taken from the US Energy Information Administration (EIA, 2018). To ensure that the count reflects sustained operation rather than sporadic service, I include only segments with at least four continuous quarters of operation by any carrier.

The resulting panel covers 1,861 segments over 16 quarters, yielding 29,776 observations. Carrier counts vary substantially within segments over this window: approximately four in five segments experience at least one change in the number of competitors. This period was also one of steady expansion in domestic service, a feature the estimation below absorbs through time effects.

I relate the number of competitors in a segment to the contemporaneous jet fuel price, controlling for the expansion and for seasonality through year and quarter effects, and for time-invariant segment characteristics through segment fixed effects. Let n_{st} denote the number of carriers in segment s and quarter t, and let w_t denote the national jet fuel spot price. I estimate a Poisson pseudo-maximum-likelihood specification,

$$E\left[n_{st} \mid w_t, \vartheta\right] = \exp\left(\beta w_t + \gamma_{y(t)} + \delta_{q(t)} + \kappa_s\right), \tag{A-25}$$

which models the proportional response of carrier counts to fuel prices, and a linear specification,

$$n_{st} = \beta w_t + \gamma_{y(t)} + \delta_{q(t)} + \kappa_s + \zeta_{st}, \tag{A-26}$$

which models the absolute response. Here $\gamma_{y(t)}$ and $\delta_{q(t)}$ are year and quarter effects and κ_s is a segment fixed effect. Standard errors are clustered by year-quarter. I also report specifications using the one-quarter lag of the fuel price.

Table A-10 reports the estimates. In all specifications, a higher fuel price is associated with fewer competitors, and the relationship is statistically significant at 1% once the strong expansion over the period and within-year seasonality are absorbed by the fixed effects. The magnitude, however, is small. The Poisson estimate of -0.078 implies that a \$1/gallon increase in the fuel price reduces the carrier count by roughly eight percent, while the linear estimate of -0.095 implies a reduction of about one-tenth of a carrier per dollar. Regressions on lagged spot prices show slightly smaller but still statistically significant effects at 5%.

The high-SCC optimal carbon tax of \$114/tCO_2 derived in section 7.2 corresponds to a jet fuel tax of approximately \$1.53 per gallon. Evaluated at this increment, both specifications imply a reduction well below a single competitor in a typical market. Since fuel price is national, the coefficient is identified from movements in a single time series net of the year and quarter effects. Moreover, the transitory nature of the identifying variation is also limiting, as forward-looking carriers likely respond less to temporary fuel-price swings than to a permanent tax. For these reasons, I interpret the estimate as evidence on the direction and order of magnitude of the response rather than as a basis for precise calibration. This limitation motivates the bounding approach in the next subsection: rather than predict the number of carriers that would exit under the optimal tax, I impose an exit substantially larger than these estimates suggest and assess whether even that extreme significantly changes the main results.

C.2.2 Bounding simulation

Next, I interpret the estimated structural model as a steady state representation and recompute a tax increase counterfactual under a forced reduction in the number of competitors. In each market served by more than one carrier, I remove the carrier with the highest passenger-weighted average fuel use per seat across its products, which is the carrier with the greatest cost exposure to a jet fuel tax and therefore the most likely to be rendered unprofitable by it. This selection rule depends only on technical fuel-burn estimates, and not on any estimated structural objects such as fixed, entry, or exit costs. Markets served by a single carrier are left unchanged. Note that this exercise is conservative in two ways, each of which removes more competition than the reduced-form evidence implies. First, I remove an entire market competitor rather than a segment competitor. Since an airline can continue to serve a market through connecting routes even after exiting a particular segment, removing a market competitor eliminates more competition than the segment-level estimates would predict. Second, I impose the removal in every market with more than one carrier, rather than in the small share of markets the estimated semi-elasticity would suggest.

When a carrier is removed, all of its products are dropped from the choice set and the Nash-Bertrand is re-solved at a higher tax rate, recomputing demands, prices, and all welfare components. Removing products rather than reassigning them to a surviving carrier is also a more conservative treatment, as it incorporates the loss of product variety as part of the cost of exit. The resulting decline in consumer surplus therefore includes a mechanical variety component and overstates the cost attributable to market power alone. I hold the tax at its no-exit optimal level ($114/t$CO_2$, as reported in section 7.2) rather than re-optimizing under the reduced structure, since endogenizing exit as the tax rises would require the dynamic model discussed above to be infeasible. The exercise thus measures the cost of the no-exit optimal tax under an adverse structural shock, not a new constrained optimum. Consistent with the treatment of fixed costs throughout the paper, any fixed cost savings from exit fall outside the welfare accounting.

Table A-11 compares the optimal-tax counterfactual under fixed market structure with the counterfactual under forced exit. The welfare-relevant summary is the *social gross average abatement cost* (SGAAC), defined as the cumulative loss of short-run private surplus (SRPS)[A-1] divided by the cumulative reduction in emissions from adding a \$114 carbon tax,

$$SGAAC\left(\tau^{*}\right)=\frac{SRPS\left(\tau^{*}\right)-SRPS\left(0\right)}{E\left(\tau^{*}\right)-E\left(0\right)}. \tag{A-27}$$

This object differs from the social gross marginal abatement cost of section 3.3, which is a derivative evaluated at a point. Since the marginal abatement cost rises along the path to the optimal tax, the average over the entire non-marginal change exceeds the baseline marginal value; as a result, the no-exit SGAAC of roughly \$244 per ton CO_2 in Table A-11 sits above the baseline SGMAC reported in Table 6.

Forcing exit raises the average cost of abatement from approximately \$244 to \$321 per ton CO_2, an increase of about a third. The additional cost arises because the reduction in emissions is somewhat smaller while the loss of private surplus is larger. The larger surplus loss is driven primarily by consumer surplus, part of which reflects the reduced product variety noted above rather than higher markups. Despite this increase, the welfare ranking across social-cost-of-carbon scenarios is preserved: the optimal tax continues to raise welfare under a high social cost of carbon, while remaining welfare-reducing under the low and medium values. Therefore, the central qualitative conclusions of the paper are thus robust to a deliberately aggressive perturbation of market structure.

The mechanism behind the modest aggregate effect is visible in the distribution of markup changes across markets. Table A-12 reports the mean change in markups by the baseline number of competitors. Markups rise materially only in markets that begin with two to four competitors, where the loss of a carrier substantially softens competitive effects, and are nearly unchanged in markets with many carriers. Note, however, that the 7-competitor case breaks the otherwise monotone decline, which could represent a composition effect from a smaller number of markets.

The markets in which exit most affects markups are precisely those that carry little aggregate weight, as markets with few competitors are overwhelmingly short-haul and account for a small

[A-1]This corresponds to the unweighted sum of consumer surplus, firms operating profits, and tax revenues.

share of passengers and emissions. The aggregate effect of forced exit is therefore contained even though its local effect on markups can be large. This result should be read as a statement about magnitude at the relevant market mix, not as evidence that market structure is unimportant. The direction of the effect confirms the reasoning in the conclusion: exit reduces competitive pressure, raises survivors' markups, and widens the markup wedge, so accounting for tax-induced exit raises the cost of abatement.

C.3 Collusive behavior

The counterfactual analyses in the paper assume that pricing is static and non-cooperative, following the Nash-Bertrand framework standard in the empirical literature on this sector. A body of evidence, however, indicates that conduct in the US airline industry can depart from this benchmark. Ciliberto et al. (2019) document pricing patterns consistent with coordination, showing that multimarket contact between carriers is associated with smaller and more rigid price gaps. Ciliberto and Williams (2014) estimate a conduct parametrization based on multimarket contact and find that implied marginal costs fall between those recovered under Nash-Bertrand and those under full coordination. If observed prices embed a degree of coordination, the markups recovered in Section 6 understate true markups, and the markup wedge in the welfare analysis is underestimated. This appendix gauges the quantitative importance of this channel.

Identification of conduct requires exclusion restrictions that the single-year cross-section used here does not support, so estimating a conduct parameter is beyond what this paper can credibly deliver. Instead, I bound the consequences of coordination by re-estimating the model at the opposite endpoint of the conduct space: all carriers serving a market jointly price all products in it, as a single multiproduct supplier would. Together with the Nash-Bertrand baseline, this endpoint brackets every intermediate degree of within-market coordination.

Implementing the coordinated endpoint involves replacing the common-ownership indicator matrix in the equilibrium system (9) with a matrix of ones, so that each first-order condition internalizes cross-price effects on every product in the market rather than only on the products of the same airline. Implied marginal costs follow from inverting these first-order conditions, and those costs enter the supply-side moments. Thus, the change in conduct requires re-estimating the joint demand-supply system of Section 5.3.

Table A-13 reports the re-estimated parameters. The demand-side moments do not involve firm behavior, and the supply equation uses its own covariates as instruments, so the demand parameters are numerically identical to the baseline estimates of Table A-2 and predicted revenues coincide with the baseline predictions; conduct operates through the implied cost vector and the cost equation (12). On the supply side, the fuel-expenditure coefficient falls from 0.745 to 0.702, consistent with the lower cost levels that coordinated pricing implies. The objective function and the Hansen J statistic are essentially unchanged as well, which illustrates the identification problem directly: both conduct assumptions fit the estimating moments equally well, so conduct cannot be recovered from fit alone.

Based on the re-estimated system, I recompute the marginal welfare analysis of Section 7.1 at the baseline tax. Table A-14 compares the marginal results under the two conduct assumptions. Under full coordination, the baseline SGMAC rises from \$196 to \$260 per ton of CO_2, an increase of about a third. The decomposition shows that this increase is due entirely to the markup wedge, which rises from \$139 to \$203, while the sales tax wedge is unchanged. This pattern follows directly from the theory: conduct determines equilibrium markups but leaves the tax schedule untouched, so it operates exclusively through the markup channel. The marginal responses of prices, passengers, and emissions are close to the baseline values, indicating that the two models disagree about the level of markups far more than about marginal behavior near the observed equilibrium.

The main qualitative conclusions are robust to both endpoints of the conduct space. Under the low SCC, the marginal welfare change is negative in both cases, so no positive carbon tax improves welfare. Under the high SCC, the marginal gain remains positive, though smaller, and re-solving the non-marginal problem of Section 7.2 under full coordination on a coarser tax grid places the optimal tax at roughly \$40/tCO_2, against \$114 in the baseline—in both cases, far below the marginal damage of \$300.

Three pieces of evidence indicate that full coordination overstates actual conduct, so that the results above are a conservative bound rather than a description of the industry. First, the out-of-sample comparison with reported financials provides information about the conduct space. Table A-15 reproduces the validation exercise of Table 5 under the collusive model. The sector-average markup reported in Form 41 financial data, 3.46 cents per available seat mile, lies between the Nash-Bertrand prediction of 2.91 and the full-coordination prediction of 4.22, roughly 40% of the distance from the former. Carrier-level deviations go in both directions, consistent with the measurement issues documented in Section 6.3 (network selection and revenue unbundling), so this comparison indicates an intermediate degree of coordination rather than as an estimate of it. Nevertheless, the position within the interval is consistent with Ciliberto and Williams (2014), whose multimarket-contact parametrization delivers marginal costs closer to the Nash-Bertrand endpoint.

Second, the full-coordination assumption fails to reproduce credible costs the low-cost carriers, a group for which collusive behavior is least expected due to their business model. For Other LCCs, the coordinated model can only rationalize observed fares by assigning near-zero marginal costs: 0.32 cents per available seat mile, against 6.69 reported. The baseline model already underpredicts costs for this group because of revenue unbundling, but the collusive model pushes the implied costs to implausible levels.

Third, the two conduct assumptions make divergent predictions about how markups relate to market structure. Under full coordination, all products in a market are priced jointly, so the number of carriers is irrelevant to pricing conditional on market characteristics. Under Nash-Bertrand, markups fall as competition intensifies. The data favors the more competitive benchmark: within haul groups, fares per mile decline with the number of carriers in short- and medium-haul markets, which jointly account for over 90% of passengers.[A-2] In contrast, the coordinated model can only rationalize this gradient by implying marginal costs that fall with the number of competitors.

To further compare the two models, let the coordination premium be the change in markups from the baseline to the coordinated model, weighted by baseline traffic. Interpreting comparisons across carrier counts, however, requires care, as the number of carriers is strongly correlated with market distance; markets served by many carriers are predominantly long-haul, so unconditional differences by carrier count mix competition and haul length. For this reason, Table A-16 reports the coordination premium by haul group and number of carriers normalized by the distance between market endpoints, alongside the mean market distance of each cell. It shows that premium is exactly zero in monopoly markets, which is expected since the two conduct assumptions coincide by construction. For all other rows, the premium is positive, quantifying how much of the markup level corresponds to conduct in each market type.

Table A-16 shows that within haul groups, the per-mile premium rises with the number of carriers. This is clearer in the short- and long-haul groups, where market distance is more uniform across carrier counts: from 6.2 to 15.7 cents per passenger-mile in short-haul markets, and from

[A-2]Table 7 reports the corresponding baseline Nash-Bertrand markup levels by haul group and number of carriers, together with each cell's traffic.

1.4 to 1.8 in long-haul ones. Therefore, the two conduct assumptions diverge precisely where competition is most intense. These short- and long-haul cells, however, account for a small share of traffic (see the counts in Table 7), so this comparison is indicative rather than conclusive; the bounding role of the coordinated model does not depend on it. As such, Nash-Bertrand remains the description of cross-market heterogeneity that the policy analysis of Section 7 relies on, with full coordination serving as a bound on levels.

Taken together, these results indicate that departures from Nash-Bertrand conduct raise the cost of abatement, which is in line with theory predictions since coordination increases the markup wedge. But the main qualitative conclusions are robust to assuming even an extreme level of coordination. The direction of the bias is reassuring for the paper's central finding: to the extent that actual conduct lies between the two endpoints, the second-best carbon tax falls even further below the Pigouvian prescription than the baseline estimates indicate. Note that, similarly to the market-structure bound of Appendix C.2, this exercise changes only one assumption at a time.

D Additional details on the data set

This appendix complements section 4 and presents additional information on definitions and variables that are calculated using multiple sources.

Markets geography. Figure A-5 shows the distribution of traffic across different markets. Lines connecting locations indicate the total number of passengers flying round trips, with thicker lines indicating a higher number of passengers. These lines connect only the endpoints of a market and are not a representation of the actual routes (i.e., they do not show connections). This map highlights the fact that the largest markets connect dense urban areas on the same coast (such as Los Angeles–San Francisco and New York–Miami), while a few are coast-to-coast or from a coast to a large city in the middle of the country (such as Chicago, Denver, and Houston).

Airline groups. When assigning ownership of flights to specific airlines, I group firms based on the controlling airline. In doing so, regional carriers that are owned by or have exclusive service agreements with a major airline are identified as being part of that major airline. This procedure is similar to Aguirregabiria and Ho (2012). Table A-17 shows how individual carriers are organized into airline groups. The last group contains five regional carriers that operate under the name of at least two major carriers, so their service contracts are not exclusive at the national level. In these cases, the major carrier is the ticketing carrier observed in the data; I use this information to assign a flight to its respective major airline group.

Jet fuel use. There are four steps in calculating the average fuel use per available seat in a route. First, I use Table T-100 in the Form 41 Traffic Database (BTS, 2018) to obtain the total departures performed and available seats by aircraft model, airline, and segment. This information allows me to calculate capacity-based use shares of each aircraft model. Second, I calculate the total jet fuel burn for each aircraft model and stage length (segment distance) using the ICAO Fuel Consumption Table in Appendix C of the ICAO Carbon Emissions Calculator Methodology v.11 (ICAO, 2018). Matching these data requires a few conversions. Table T-100 identifies aircraft models with DOT codes instead of the standard IATA Designator Codes used in ICAO's table; I manually match the 44 aircraft models in (the subset of) T-100 to IATA Designator Codes. Moreover, ICAO's table provides total fuel burn in kilograms for stage lengths in nautical miles (nm). I convert jet fuel mass to volume using the reference density of 820 kg/m^3 at 15 °C—approximately 3.10 kg/gallon at 60 °F—for Jet Fuel A, the standard used in US aviation (NREL, 2001). Also, one nautical mile corresponds to 1.15078 miles. Total fuel burn is tabulated at values of 125, 250, 500, 750, 1000, and subsequent increments of 500 nm up to 8500 nm, depending on the range of the model. I use linear interpolation to calculate fuel burn between tabulated stage lengths.

Figure A-6 illustrates the importance of factoring stage lengths instead of relying on average fuel efficiency by aircraft model. This figure shows how fuel burn per available seat-mile decreases with increased distance flown for the eight most used models.[A-3] The "hockey stick" shape of efficiency curves represents the fact that taking off and climbing consume a substantial fraction of the fuel allocated to a flight. Since the cruising stage is more fuel efficient, longer flights burn less fuel per mile on average. ICAO's fuel burn data are representative of the average conditions in which these aircraft models are used in industry. Actual fuel consumption in a given trip—information not publicly available—varies by weight load, flight plan, weather conditions, and other factors. Nevertheless, passenger occupation rate, or load factor, has a small effect on fuel burn, as passengers and baggage typically account for only 15% of the take-off weight (Borenstein & Rose, 2014).

[A-3] The number of available seats of an aircraft model varies by plane, with some airlines using more dense configurations. The calculation in this graph uses the average number of seats across all planes.

Third, I calculate the average fuel use for each segment, airline, and quarter. The fuel use of an aircraft model in a segment is the total fuel burn, calculated in the previous step, divided by the average number of seats available for that model and airline. I allow the number of seats to vary by airline in order to accommodate different seat configurations. Then, for each segment and airline in a quarter, I calculate average fuel efficiency across aircraft models used (if the airline used more than one), weighting each model by the capacity-based use share in that quarter.

Fourth and finally, I calculate the average fuel use per available seat in a product (i.e., by route, airline, and quarter) by summing the average fuel use of each segment in a route for that airline and quarter. Fuel efficiency varies substantially at every travel distance, and the mode of its distribution shifts toward higher efficiency on longer routes (see Figure A-3).

Jet fuel expenditure. Jet fuel prices are gathered from the US Energy Information Administration (EIA, 2018). I use average end-user sales prices by quarter and region. Regions are based on the Petroleum Administration for Defense Districts (PADD). The data cover all five districts (and sub-districts): West Coast, Rocky Mountain, Gulf Coast, Midwest, and East Coast. The East Coast PADD has three sub-districts: New England, Central Atlantic, and Lower Atlantic. End-user sales prices include delivery costs, which capture relevant variations in fuel prices across the US. Despite this spatial variation, end-user sales prices closely track changes in spot prices, as illustrated in Figure A-7, panel (a). Panel (b) in this Figure shows the variation of end-user sales prices across regions and compares it with spot prices.

To calculate jet fuel expenditure per available seat, I multiply the average jet fuel use in each segment by its respective jet fuel price. The price assigned to each segment corresponds to the PADD where the departing airport is located. Then, for each product, I sum fuel expenditures of all the segments of a product's route.

Load factors. Fuel use, costs, and emission intensities in this paper are measured per available seat, while quantities count passengers, so the analysis holds the load factor connecting these two units fixed, as discussed in Section 5.3. Figure A-8 reports monthly load factors for the six largest US carriers over 2014–2019, calculated as revenue passenger-miles divided by available seat-miles in scheduled domestic service using Schedule T-1 of the Form 41 Traffic Database (BTS, 2018). Two features of these series support the assumption. First, occupancy is stable at the annual frequency: the traffic-weighted load factor of these carriers lies between 85.1 and 85.8 percent in every year of the period. Second, almost all of the movement within a year follows the calendar rather than operating conditions. For instance, regressing monthly carrier load factors on calendar-month indicators shows that the month of the year alone accounts for roughly two thirds of the observed variation, while year indicators account for less than one percent of it.

The stability of load factors in this period is informative because operating conditions were not stable. Jet fuel prices fell by roughly two thirds between 2014 and 2016 and then recovered by 2019, as panel (a) of Figure A-7 shows. However, load factors did not follow these movements. This shows that traffic tracks capacity rather than occupancy: twelve-month changes in available seat-miles and in revenue passenger-miles have nearly identical dispersion across these carriers and a correlation of 0.98. Room for occupancy to rise is also limited, since carriers manage load toward a target rather than a full aircraft, as the last seats sold raise the cost of accommodating disrupted itineraries. The highest monthly load factor recorded by any of these carriers over this period is 91.8 percent, and operating every month of 2018 at each carrier's own best month while carrying the same traffic would have required about 6 percent fewer available seat-miles.

References

Aguirregabiria, V., & Ho, C.-Y. (2012). A dynamic oligopoly game of the US airline industry: Estimation and policy experiments. *Journal of Econometrics*, *168*(1), 156–173.

Belobaba, P., Odoni, A., & Barnhart, C. (2015). *The global airline industry.* John Wiley & Sons.

Borenstein, S., & Rose, N. L. (2014). How airline markets work... or do they? regulatory reform in the airline industry. In *Economic regulation and its reform: What have we learned?* (pp. 63–135). University of Chicago Press.

BTS. (2018). *Data products: Transtats.* Bureau of Transportation Statistics. Retrieved June, 2020, from `https://www.transtats.bts.gov`

Ciliberto, F., Watkins, E., & Williams, J. W. (2019). Collusive pricing patterns in the us airline industry. *International Journal of Industrial Organization*, *62*, 136–157.

Ciliberto, F., & Williams, J. W. (2014). Does multimarket contact facilitate tacit collusion? Inference on conduct parameters in the airline industry. *The RAND Journal of Economics*, *45*(4), 764–791.

Conlon, C., & Gortmaker, J. (2020). Best practices for differentiated products demand estimation with PyBLP. *The RAND Journal of Economics*, *51*(4), 1108–1161.

Dekle, R., Eaton, J., & Kortum, S. (2008). *Global rebalancing with gravity: Measuring the burden of adjustment* (Working Paper No. 13846). National Bureau of Economic Research. Retrieved from `https://www.nber.org/papers/w13846` doi: 10.3386/w13846

EIA. (2018). *Refiner Petroleum Product Prices by Sales Type.* U.S. Energy Information Administration. Retrieved June, 2020, from `https://www.eia.gov/dnav/pet/pet_pri_refoth_dcu_nus_m.htm`

EPA. (2018). *Greenhouse gas inventory guidance: Indirect emissions from events and conferences.* U.S. Environmental Protection Agency. Retrieved from `https://www.epa.gov/sites/`

default/files/2018-12/documents/indirectemissions_draft2_12212018_b_508pass_3.pdf

ICAO. (2018). *Carbon Emissions Calculator Methodology, Version 11.* International Civil Aviation Organization. Retrieved June, 2020, from https://www.icao.int/environmental-protection/CarbonOffset/Documents/Methodology%20ICAO%20Carbon%20Calculator_v11-2018.pdf

NREL. (2001). *Investigation of Byproduct Application to Jet Fuel.* National Renewable Energy Laboratory. Retrieved June, 2020, from https://www.nrel.gov/docs/fy02osti/30611.pdf

Sanderson, E., & Windmeijer, F. (2016). A weak instrument F-test in linear IV models with multiple endogenous variables. *Journal of Econometrics*, *190*(2), 212–221.

Figure A-1: Welfare consequences of different levels of carbon taxation.

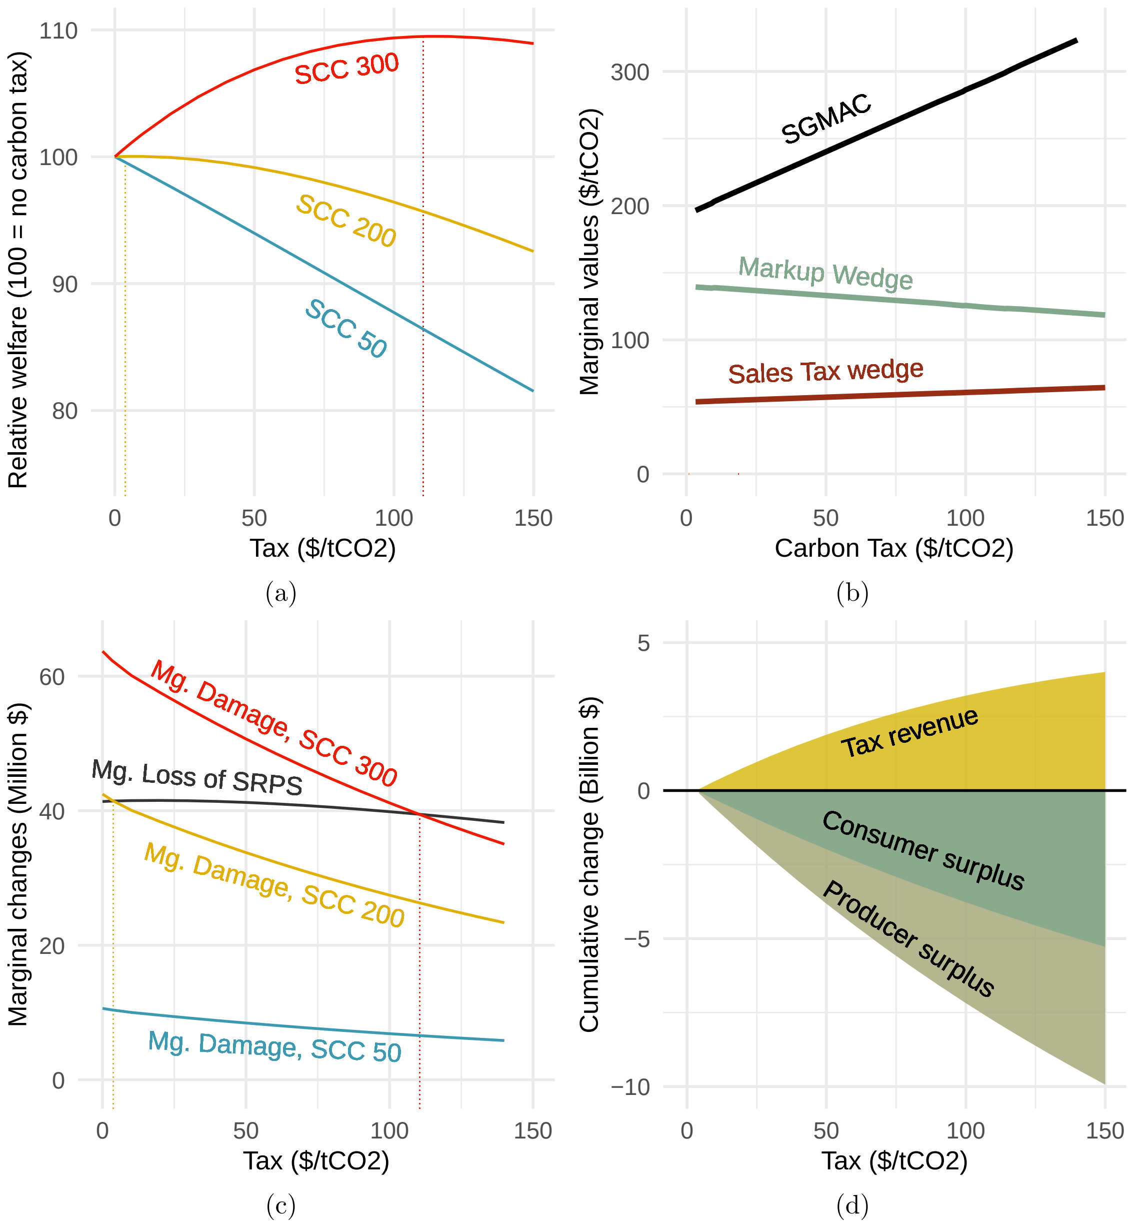

Notes: *Mg. Damage* corresponds to the marginal benefit of lowering emissions, which depends on the social cost of carbon (SCC) and the emission reductions at each tax level. *Loss of SRPS* (short-run private surplus) is analogous to the aggregate private cost of reducing emissions with a carbon tax.

Figure A-2: Baseline characteristics of flights by carbon intensity quintiles.

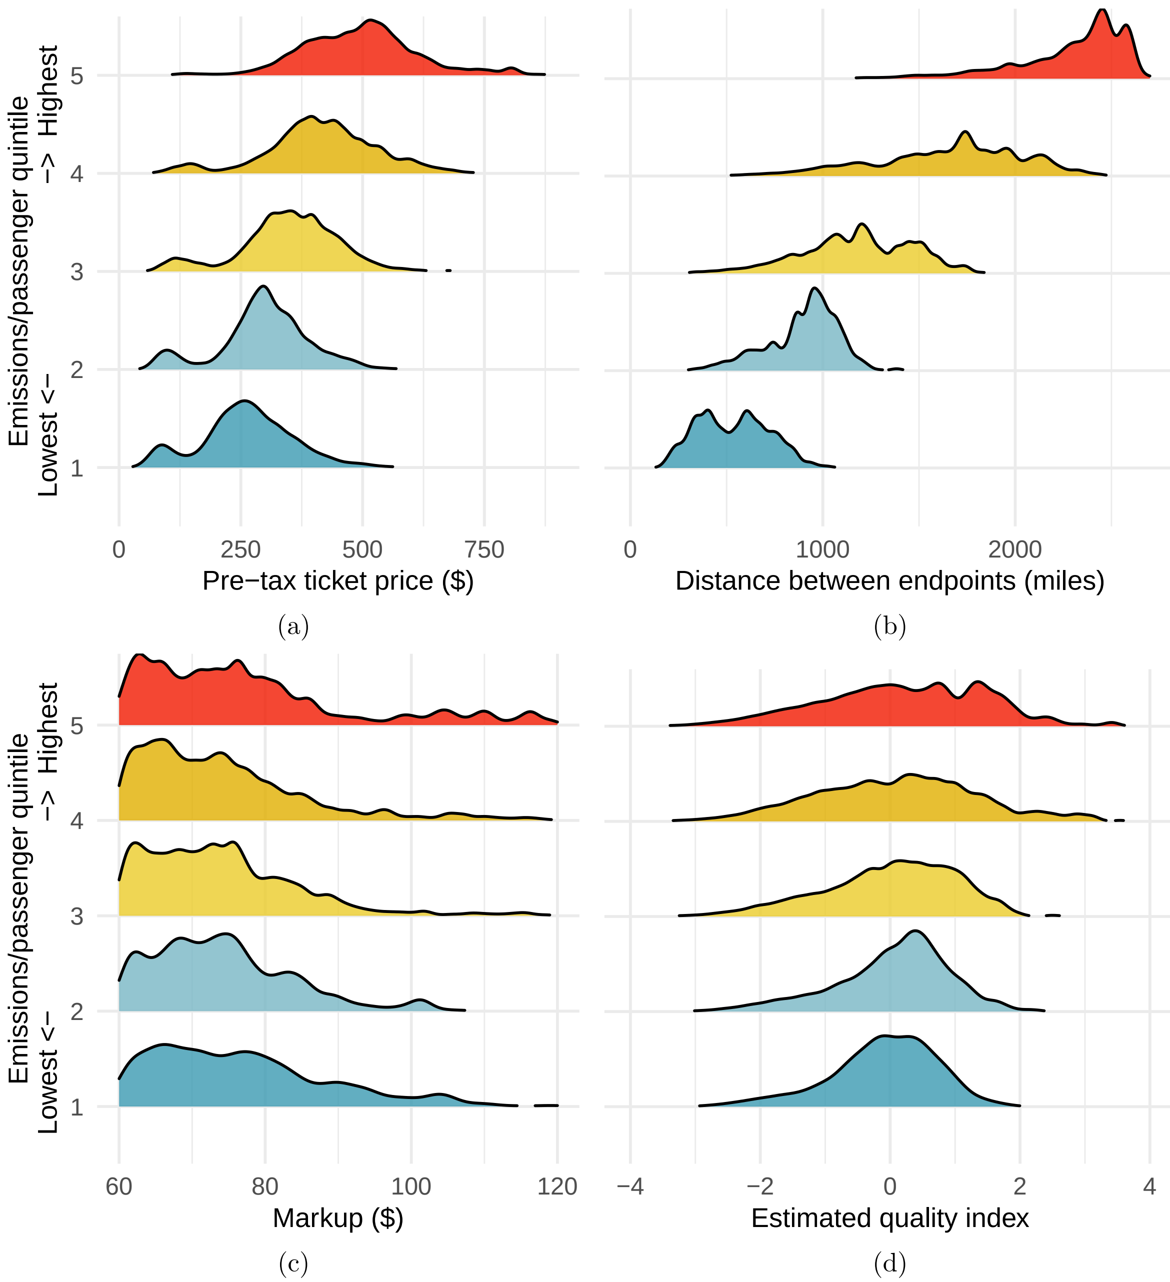

Notes: Curves show kernel densities for each carbon intensity quintile. Distributions are weighted by total emissions, so each quintile has approximately the same baseline aggregate emissions. Quintile cut-offs are approximately 320, 410, 530, and 710 Kg of CO_2 per passenger.

Figure A-3: Fuel efficiency and travel distance.

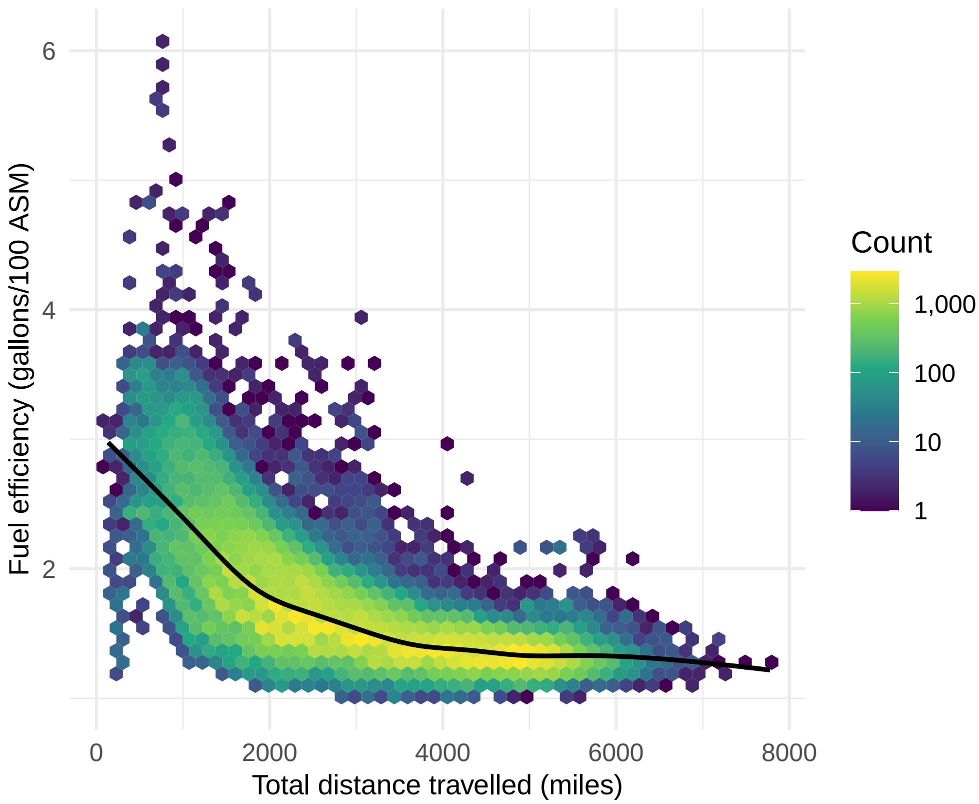

Notes: Joint distribution of fuel efficiency (gallons per hundred available seat-miles) and total travel distance across products. Color denotes the number of products (flight-quarters) in each distance-efficiency cell, on a log scale. The black line shows the local average (GAM) trend.

Figure A-4: Passenger traffic in cities and metro areas included in the data set.

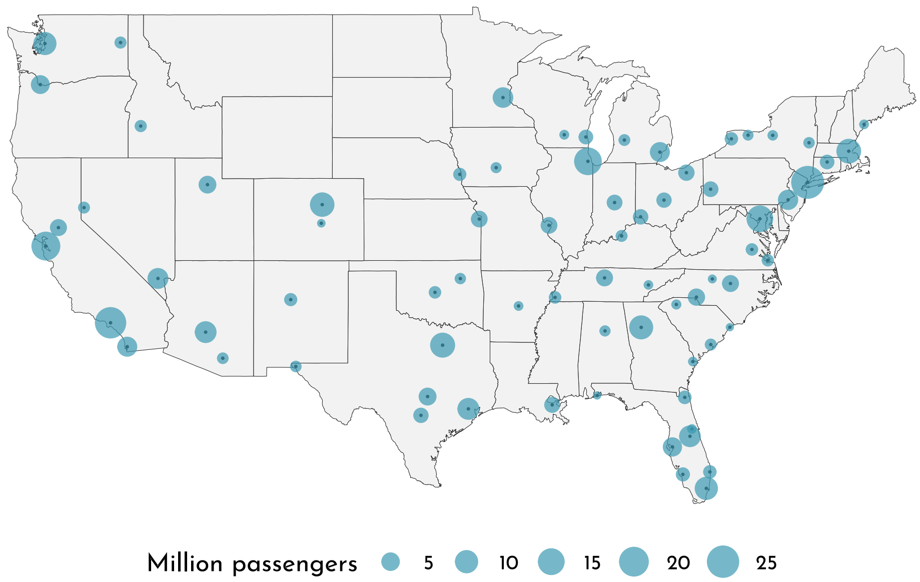

Source of data: US Bureau of Transportation Statistics. Notes: traffic is measured in passengers enplaned in domestic flights from all airports within a city or metro area.

Figure A-5: Passenger traffic between cities and metro areas.

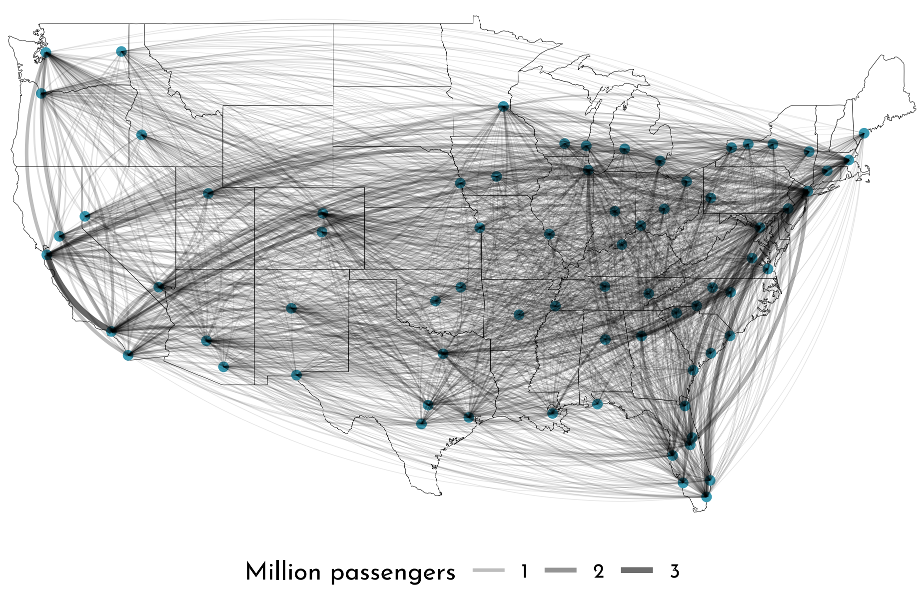

.

Notes: each segment connects only the endpoints (city or metro area) of a round trip, regardless of any connections in between. Traffic measures the number of passengers enplaned in round-trip flights between any airport at the endpoints in either direction.

Figure A-6: Fuel efficiency by stage length for the aircraft models most used.

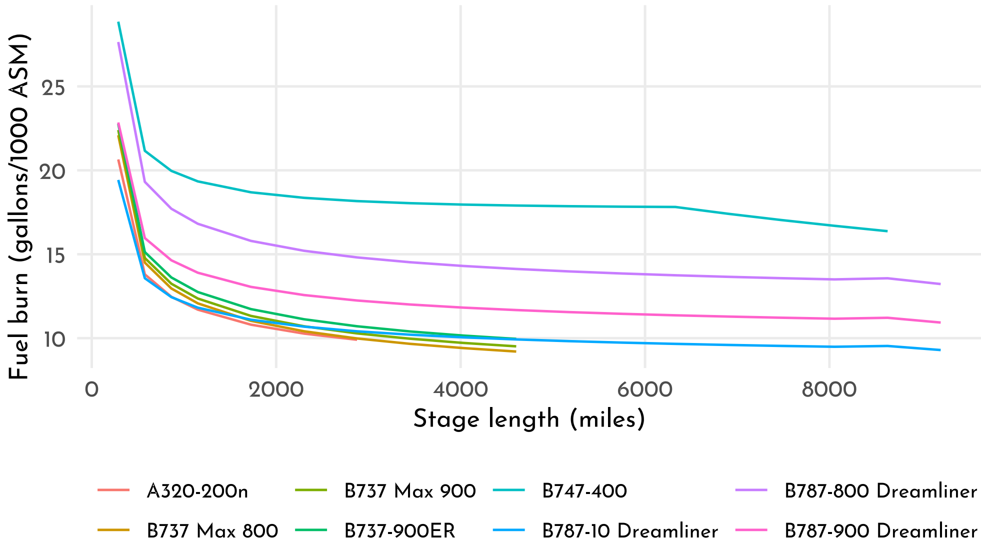

Notes: Fuel efficiency is measured in gallons of jet fuel per 1000 available seat-miles. The eight models shown in this graph correspond to those that offered the highest aggregate capacity (in available seat-miles) in 2018.

Figure A-7: Jet fuel and crude oil prices.

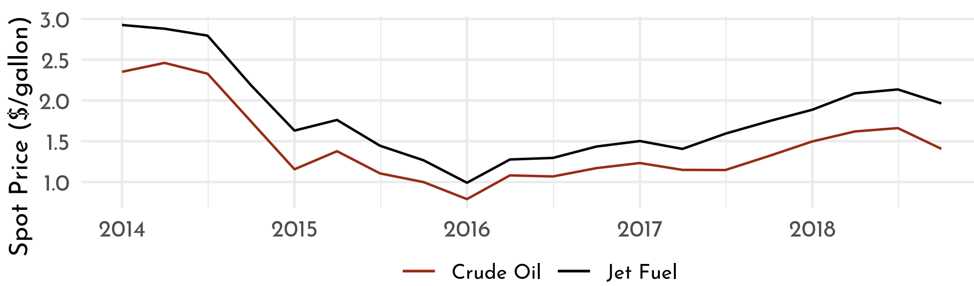

(a) Jet fuel and crude oil spot prices.

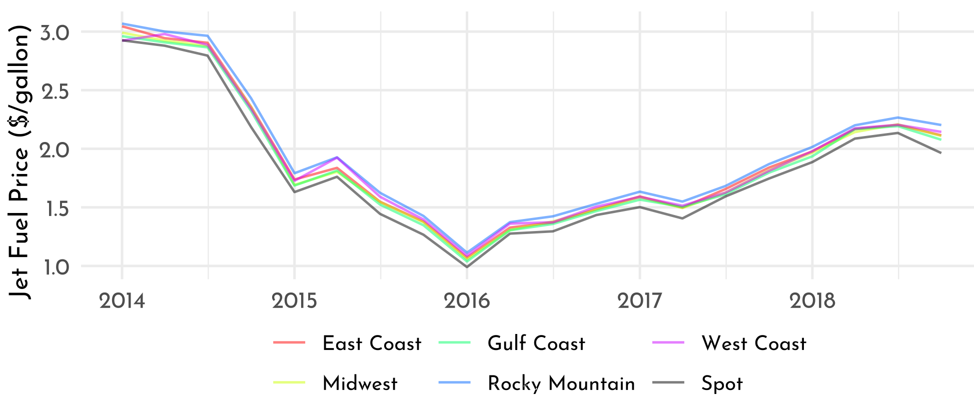

(b) Jet fuel spot and regional end-user sales prices.

Notes: Spot prices are for US Gulf Coast Kerosene-Type Jet Fuel and Crude Oil WTI. Regional (PADD) prices are average end-user sales prices by quarter. *Source: US Energy Information Administration.*

Figure A-8: Monthly load factors of the six largest US carriers.

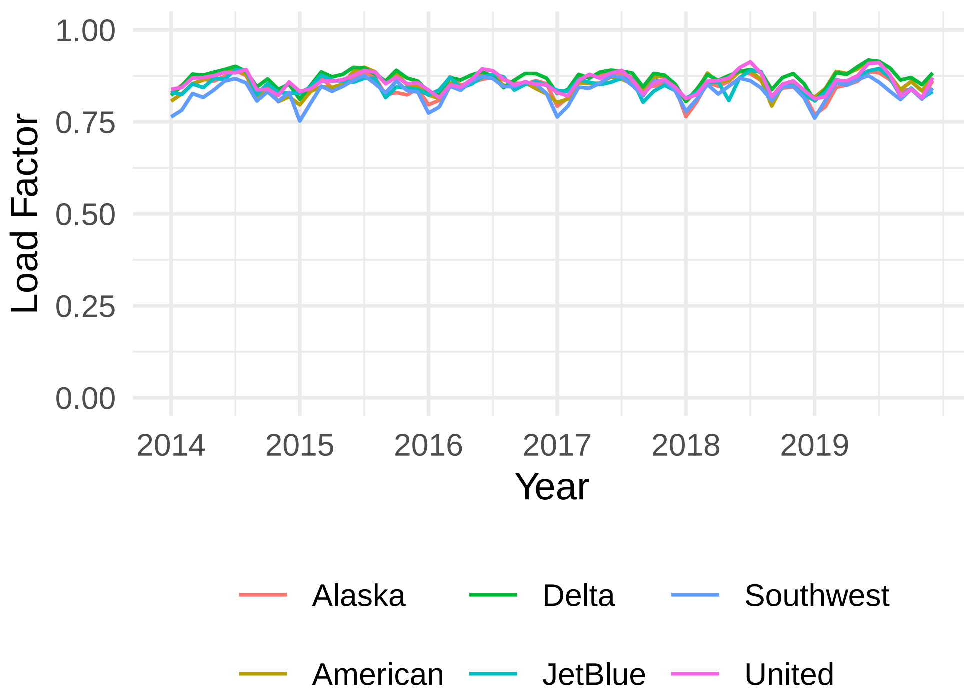

Notes: The load factor is revenue passenger-miles divided by available seat-miles. Series cover scheduled domestic passenger service and are calculated from Schedule T-1 of the Form 41 Traffic Database. *Source: US Bureau of Transportation Statistics.*

Table A-1: Results for the estimation of sufficient statistics.

		Cost shifter only		All instruments	
	OLS	2SLS	1st stage	2SLS	1st stage
	ln(Passengers)	ln(Passengers)	ln(Price)	ln(Passengers)	ln(Price)
ln(Price) [ε]	−0.472	−1.848		−1.797	
	(0.055)	(0.138)		(0.136)	
Products in market	0.015	0.015	0.0004	0.015	0.0004
	(0.001)	(0.001)	(0.0001)	(0.001)	(0.0001)
% of nonstop flights	−0.059	−0.254	−0.060	−0.246	−0.054
	(0.146)	(0.207)	(0.061)	(0.205)	(0.061)
Excluded instruments					
log(Fuel expenditure)			0.756		0.751
			(0.043)		(0.043)
Potential LCC entrants					0.009
					(0.003)
Potential legacy entrants					−0.012
					(0.007)
Fixed effects					
City pair	Yes	Yes	Yes	Yes	Yes
Quarter	Yes	Yes	Yes	Yes	Yes
Observations	19,750	19,750	19,750	19,750	19,750
R^2	0.963	0.959	0.921	0.959	0.921
R^2 (within)	0.205	0.113	0.115	0.120	0.117
Conditional F-statistic			105.0		67.8

Notes: Standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The conditional F-statistic is calculated following Sanderson and Windmeijer (2016). The within R^2 is calculated after residualizing on the fixed effects.

Table A-2: Results for the joint demand-supply estimation of the structural model.

Demand	$ln(s_{kt}/s_{0t})$	*Supply*	c_{kt}
Price (\$100) [$-\boldsymbol{\alpha_0}$]	-0.936	Fuel expenses [$\boldsymbol{\rho}$]	0.745
	(0.116)		(0.003)
Price × Demeaned income [$-\boldsymbol{\alpha_1}$]	0.100	Total ramp-to-ramp time (h)	0.147
	(0.054)		(0.137)
ln(share within nest) [$\mathbf{1-\boldsymbol{\lambda}}$]	0.443	× *American*	-0.041
	(0.049)		(0.015)
Departures per week	0.037	× *Delta*	0.047
	(0.002)		(0.007)
Number of stops	-0.734	× *United*	-0.018
	(0.071)		(0.008)
Market distance (100 mi.)	0.103	× *Alaska*	-0.020
	(0.017)		(0.008)
Market distance squared (100 mi.)^2	-0.055	× *JetBlue*	-0.017
	(0.041)		(0.032)
Connection extra distance (100 mi.)	-0.079	× *Other low-cost*	-0.126
	(0.009)		(0.010)
Connection extra distance squared (100 mi.)^2	0.300		
	(0.053)		
Share of delayed departures	-0.560		
	(0.110)		
Destinations from origin	0.014		
	(0.003)		
Observations	267,967		
Objective function minimum	1.33 $\times 10^{-5}$		
Hansen J-stat	3.56 (p = 0.31)		

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The demand equation includes fixed effects for airline, origin airport-by-quarter, and destination airport-by-quarter. The supply equation includes fixed effects for quarter and for each route airport-by-airline. Southwest is the base (omitted) airline in the interaction with ramp-to-ramp time. Baseline estimates are reported in Table A-2; the demand columns coincide because the demand-side moments do not involve conduct.

Table A-3: Estimates of demand parameters using 2SLS.

	$\ln(s_k/s_0)$
Price ($100) [$-\alpha_0$]	-0.880
	(0.114)
Price x Demeaned income [$-\alpha_1$]	0.110
	(0.054)
ln(share within nest) [1-λ]	0.436
	(0.049)
Departures per week	0.036
	(0.002)
Number of stops	-0.743
	(0.070)
Market distance (100 mi.)	0.099
	(0.016)
Market distance squared	-0.001
	(0.0004)
Connection extra distance (100 mi.)	-0.081
	(0.008)
Connection extra distance squared	0.003
	(0.001)
Share of delayed departures	-0.592
	(0.099)
Destinations from origin	0.012
	(0.002)
Fixed effects	
Airline	Yes
Origin-by-quarter	Yes
Destination-by-quarter	Yes
Observations	267,967

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together).

Table A-4: First stage regressions of the demand model using 2SLS.

	Dependent variable:			
	Price	Price × income	ln(share within nest)	Departures per week
Number of stops	−0.205	0.036	−1.179	−0.030
	(0.022)	(0.068)	(0.022)	(0.206)
Market distance (100 mi.)	−0.185	0.068	0.016	−0.970
	(0.010)	(0.036)	(0.013)	(0.084)
Market distance squared	0.005	0.001	0.004	0.031
	(0.0002)	(0.001)	(0.0004)	(0.002)
Connection extra distance	−0.009	−0.024	−0.152	−0.920
	(0.005)	(0.011)	(0.005)	(0.053)
Connection extra distance squared	−0.001	0.001	0.010	0.076
	(0.0003)	(0.001)	(0.0004)	(0.004)
Share of delayed departures	−0.154	−0.383	−0.416	−6.059
	(0.094)	(0.258)	(0.107)	(0.798)
Destinations from origin	0.011	0.018	−0.010	−0.086
	(0.001)	(0.002)	(0.001)	(0.006)
Airlines in market	0.211	−0.060	−0.129	−0.629
	(0.015)	(0.047)	(0.022)	(0.144)
Rivals' products in market	0.003	−0.0005	−0.002	0.008
	(0.0002)	(0.001)	(0.0004)	(0.002)
Rivals' % of nonstop flights	−0.512	−0.602	−0.222	7.351
	(0.077)	(0.552)	(0.225)	(1.266)
Potential legacy entrants	0.383	0.074	−0.049	0.006
	(0.031)	(0.119)	(0.062)	(0.428)
Potential LCC entrants	0.171	−0.029	0.058	−0.857
	(0.012)	(0.039)	(0.017)	(0.113)
Fuel expenditure	0.013	0.003	−0.015	−0.066
	(0.001)	(0.004)	(0.001)	(0.010)
Compl. segment density	0.003	0.007	0.007	0.320
	(0.0004)	(0.001)	(0.0003)	(0.005)
Fixed effects				
Airline	Yes	Yes	Yes	Yes
Origin airport-by-quarter	Yes	Yes	Yes	Yes
Destination airport-by-quarter	Yes	Yes	Yes	Yes
Observations	267,967	267,967	267,967	267,967
F-statistic	316.76	3216.11	956.98	352.8
Conditional F-statistic	27.15	12.37	28.01	40.0

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The conditional F-statistic is calculated following Sanderson and Windmeijer (2016).

Table A-5: Estimates of supply parameters using 2SLS.

	$\widehat{c}_k$
Fuel expenditure/avail. seat [ρ]	0.672
	(0.162)
Total ramp-to-ramp time (h)	0.157
	(0.017)
$\times$ *American*	−0.040
	(0.008)
$\times$ *Delta*	0.047
	(0.008)
$\times$ *United*	−0.017
	(0.008)
$\times$ *Alaska*	−0.021
	(0.033)
$\times$ *JetBlue*	−0.015
	(0.010)
$\times$ *Other low-cost*	−0.125
	(0.007)
Fixed effects	
Quarter	Yes
Each route airport-by-airline	Yes
Observations	267,967

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together).

Table A-6: Estimates of demand parameters under alternative specifications.

	Dependent variable:				
	$\ln(s_k/s_0)$				
	(1)	(2)	(3)	(4)	(5)
	OLS	2SLS	2SLS	2SLS	2SLS
		No FE	No income inter.	Airline nests	Flight type nests
Price (\$100) $[-\alpha]$	0.008	−0.252	−0.800	0.087	−0.158
	(0.0004)	(0.157)	(0.110)	(0.086)	(0.025)
Price x Demeaned income $[-\alpha_1]$	−0.042	−0.041		−0.018	−0.037
	(0.002)	(0.019)		(0.026)	(0.010)
ln(share within nest) [1-λ]	0.009	0.353	0.366	0.399	1.034
	(0.002)	(0.059)	(0.050)	(0.036)	(0.027)
Departures per week	0.791	0.018	0.036	0.029	0.0003
	(0.008)	(0.002)	(0.002)	(0.002)	(0.001)
Number of stops	−0.344	−0.781	−0.848	−1.193	−0.352
	(0.013)	(0.065)	(0.073)	(0.018)	(0.028)
Market distance (100 mi.)	−0.006	−0.009	0.077	0.007	−0.001
	(0.009)	(0.032)	(0.015)	(0.009)	(0.003)
Market distance squared	−0.002	0.001	0.0001	−0.001	0.001
	(0.0004)	(0.001)	(0.0003)	(0.0003)	(0.0001)
Connection extra distance (100 mi.)	−0.027	−0.096	−0.089	−0.073	−0.013
	(0.002)	(0.014)	(0.009)	(0.007)	(0.004)
Connection extra distance squared	0.001	0.005	0.004	0.002	0.001
	(0.0001)	(0.001)	(0.001)	(0.001)	(0.0002)
Share of delayed departures	−0.378	−0.726	−0.603	−0.681	0.143
	(0.054)	(0.095)	(0.094)	(0.074)	(0.033)
Destinations from origin	−0.001	−0.003	0.011	0.009	0.0005
	(0.0002)	(0.002)	(0.002)	(0.002)	(0.001)
Fixed effects					
Airline	Yes	Yes	Yes	Yes	Yes
Origin-by-quarter	Yes	No	Yes	Yes	Yes
Destination-by-quarter	Yes	No	Yes	Yes	Yes
Nesting	All flights	All flights	All flights	By airline	Stop vs. nonstop
Observations	267,967	267,967	267,967	267,967	267,967

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together).

Table A-7: Summary statistics for the estimation of the elasticity of substitution.

	Mean	St. Dev.	Min	Median	Max
Fuel use (gallons/100 ASM)	1.716	0.519	0.897	1.590	4.958
Fuel spot price ($/gallon)	1.787	0.529	0.991	1.630	2.925
Observations			1,634		
Time periods (quarter)			20		
Airlines			14		
Aircraft models			43		

Table A-8: Estimates of the elasticity of substitution between fuel and non-fuel inputs.

	Dependent variable:		
	Fuel use	log(Fuel use)	
	(1)	(2)	(3)
	Linear	Log-log	CES implied
Spot price	−0.013		
	(0.008)		
log(Spot price)		−0.013	
		(0.007)	
log(Spot price) × (1 - Fuel cost share)			−0.060
			(0.012)
Implied elasticity [σ]	0.008	0.020	0.060
Fixed effects			
Aircraft model	Yes	Yes	Yes
Airline	Yes	Yes	Yes
Observations	1,634	1,634	1,470
R^2	0.909	0.918	0.916

Notes: Standard errors are shown in parentheses. Fuel use is measured in gallons per hundred available seat-miles; fuel spot prices are in dollars per gallon.

Table A-9: Comparison of marginal effect estimates across methods with input substitution.

	Sufficient statistics			Structural model		
Elasticity of substitution (σ)	0	0.04	0.08	0	0.04	0.08
Marginal costs (\$/tCO$_2$)						
Markup wedge	185.12	172.45	161.39	139.34	133.16	127.50
Sales tax wedge	54.10	50.39	47.16	53.74	51.36	49.18
Baseline (equivalent) carbon tax	3.28	3.28	3.28	3.28	3.28	3.28
Social gross marginal abatement cost (SGMAC)	242.51	226.12	211.84	196.36	187.65	179.69
Marginal changes on market aggregates (per tCO$_2$)						
Mean ticket price	+0.12%	+0.12%	+0.12%	+0.09%	+0.09%	+0.09%
Passengers	-0.22%	-0.22%	-0.22%	-0.30%	-0.30%	-0.30%
Emissions	-0.24%	-0.25%	-0.27%	-0.37%	-0.38%	-0.40%

Table A-10: Effects of jet fuel prices on the presence of carriers on segments.

	Dependent variable: Carrier count			
	(1)	(2)	(3)	(4)
	Poisson	OLS	Poisson	OLS
Spot price	-0.078	-0.095		
	(0.028)	(0.031)		
Lagged spot price			-0.066	-0.079
			(0.032)	(0.043)
Fixed effects				
Year	Yes	Yes	Yes	Yes
Quarter	Yes	Yes	Yes	Yes
Segment	Yes	Yes	Yes	Yes
Observations	29,776	29,776	27,840	27,915
Pseudo R^2	0.273	–	0.269	–
R^2	–	0.856	–	0.861

Notes: Standard errors clustered by year × quarter are shown in parentheses.

Table A-11: Optimal-tax counterfactual under fixed structure and forced exit.

	Fixed structure	Forced exit
Changes on market aggregates		
Mean ticket price	+10.41%	+10.64%
Passengers [millions]	-27.60% [-37.67]	-26.38% [-34.16]
Emissions [million tons]	-32.34% [-18.61]	-30.82% [-16.60]
Private welfare changes ($ billions)		
Consumer surplus	-4.22	-4.61
Operating profit	-3.79	-3.89
Fuel tax revenue	4.25	4.06
Sales tax revenue	-0.77	-0.88
Short-Run Private Surplus	-4.53	-5.32
Social Gross Average Abatement Cost (SGAAC, $/tCO$_2$)	243.59	320.65
Net social welfare changes ($ billions)		
SCC of: $50	-3.60	-4.31
$200	-0.81	-1.26
$300	+1.05	+0.76

Notes: Changes are relative to the corresponding no-tax baseline. Optimal tax is $114/tCO$_2$ set under the high-SCC scenario.

Table A-12: Mean change in markups by baseline number of competitors.

Baseline competitors	Mean markup change ($)	Markets × quarters	Baseline passenger share (%)
1	0.00	1648	0.70
2	5.13	3485	2.39
3	1.91	4446	7.22
4	0.94	5680	19.27
5	0.46	2938	26.42
6	0.29	1150	25.34
7	0.61	323	15.31
8	0.27	69	2.27
9	0.19	11	1.09

Notes: Change in markups for monopoly is zero by construction (no exit imposed). The third column counts the number of markets times the number of quarters that markets had the corresponding number of competitors to account for changes in structure during the sample period.

Table A-13: Results for the joint demand-supply estimation under full within-market coordination.

Demand	$ln(s_{kt}/s_{0t})$	*Supply*	c_{kt}
Price (\$100) [$-\boldsymbol{\alpha_0}$]	-0.936	Fuel expenses [$\boldsymbol{\rho}$]	0.702
	(0.116)		(0.003)
Price × Demeaned income [$-\boldsymbol{\alpha_1}$]	0.100	Total ramp-to-ramp time (h)	0.142
	(0.054)		(0.136)
ln(share within nest) [$\mathbf{1-\boldsymbol{\lambda}}$]	0.443	× *American*	-0.026
	(0.049)		(0.015)
Departures per week	0.037	× *Delta*	0.061
	(0.002)		(0.007)
Number of stops	-0.734	× *United*	-0.002
	(0.071)		(0.008)
Market distance (100 mi.)	0.103	× *Alaska*	-0.012
	(0.017)		(0.008)
Market distance squared (100 mi.)^2	-0.055	× *JetBlue*	-0.003
	(0.041)		(0.032)
Connection extra distance (100 mi.)	-0.079	× *Other low-cost*	-0.112
	(0.009)		(0.010)
Connection extra distance squared (100 mi.)^2	0.300		
	(0.053)		
Share of delayed departures	-0.560		
	(0.110)		
Destinations from origin	0.014		
	(0.003)		
Observations	267,967		
Objective function minimum	1.33 $\times 10^{-5}$		
Hansen J-stat	3.56 (p = 0.31)		

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The demand equation includes fixed effects for airline, origin airport-by-quarter, and destination airport-by-quarter. The supply equation includes fixed effects for quarter and for each route airport-by-airline. Southwest is the base (omitted) airline in the interaction with ramp-to-ramp time. Baseline estimates are reported in Table A-2; the demand columns coincide because the demand-side moments do not involve conduct.

Table A-14: Marginal welfare effects at the baseline tax under alternative conduct assumptions.

	Nash-Bertrand	Full coordination
Marginal costs (\$/tCO_2)		
Markup wedge	139.34	203.04
Sales tax wedge	53.74	53.74
Baseline (equivalent) carbon tax	3.28	3.28
Social gross marginal abatement cost (SGMAC)	196.36	260.06
Marginal changes on market aggregates (per tCO_2)		
Mean ticket price	+0.09%	+0.09%
Passengers	−0.30%	−0.28%
Emissions	−0.37%	−0.34%

Notes: The full-coordination column is computed from a re-estimation of the model in which all carriers in a market jointly price all of its products.

Table A-15: Predicted costs and markups under alternative conduct assumptions versus reported financials.

	Cost (cents/ASM)			Markup (cents/ASM)		
Airline	N-B	Coord.	Reported	N-B	Coord.	Reported
Southwest	9.99	8.58	9.44	3.49	4.91	4.26
American	11.83	10.66	10.37	2.74	3.91	4.28
Delta	12.97	11.83	10.89	2.94	4.08	4.92
United	11.33	9.92	10.75	2.58	3.99	1.85
Other LCCs	1.77	0.32	6.69	2.44	3.89	1.62
JetBlue	8.06	6.60	9.46	2.67	4.13	2.62
Alaska	7.15	5.89	8.37	2.93	4.18	3.11
Sector average	10.12	8.81	9.58	2.91	4.22	3.46

Notes: N-B denotes the Nash-Bertrand baseline and Coord. the full-coordination re-estimation. Reported values are computed from Form 41 financial data as in Table 5. Predicted revenues are omitted because they are essentially identical across the two conduct assumptions because the demand-side estimates are unaffected by conduct. Rows are ordered by market share as in Table 5.

Table A-16: Mean change in markups from Nash-Bertrand to full coordination, by haul group and number of carriers.

Haul	Carriers	Passengers (M)	Δ markup/mile (¢/pax-mi)	Mean market distance (mi)
Short	1	0.17	0.0	238
	2	0.85	6.2	221
	3	1.72	11.0	238
	4	1.92	15.2	243
	5+	1.79	15.7	244
Medium	1–2	1.39	2.1	565
	3	3.80	3.8	635
	4	19.58	4.4	720
	5	28.07	3.9	897
	6	31.14	4.0	1,058
	7	29.61	4.6	1,131
	8+	8.04	3.9	1,502
Long	1–4	0.27	1.4	2,455
	5–7	1.83	1.6	2,478
	8+	6.29	1.8	2,482

Notes: Entries are passenger-weighted means of the product-level change in markups between the two conduct assumptions, with weights held at the baseline equilibrium to keep them invariant to conduct. Haul groups are based on the distance between market endpoints and follow the US Environmental Protection Agency suggested classification: short haul $\leq$ 300 miles, medium haul $>$ 300 and $\leq$ 2,300 miles (EPA, 2018). Per-mile markups are normalized by market distance. Cells pool carrier counts where markets are sparse. The change in markets with a single carrier is zero by construction, as the two conduct assumptions coincide. Carrier-count bins with small passenger counts are pooled within haul groups.

Table A-17: Airline groups and types in the data set.

Carrier name (code)	Airline group	Type of airline
American Airlines Inc. (AA) PSA Airlines Inc. (16, OH) Envoy Air (MQ)	American	Legacy
Delta Air Lines Inc. (DL) Endeavor Air Inc. (9E)	Delta	Legacy
United Air Lines Inc. (UA) ExpressJet Airlines LLC (EV) Island Air Hawaii (WP) Air Wisconsin Airlines Corp (ZW)	United	Legacy
Alaska Airlines Inc. (AS) Horizon Air (QX)	Alaska	Mixed
Southwest Airlines Co. (SW)	Southwest	Low-cost
JetBlue Airways (B6)	JetBlue	Low-cost
Frontier Airlines Inc. (F9)	Frontier	Low-cost
Allegiant Air (G4)	Allegiant	Low-cost
Spirit Air Lines (NK)	Spirit	Low-cost
Sun Country Airlines (SY)	Sun Country	Low-cost
GoJet Airlines (G7) Compass Airlines (CP) SkyWest Airlines Inc. (OO) Mesa Airlines Inc. (YV) Republic Airline (YX)	American, Delta, or United (varies by market)	Regional partners