Second-best carbon taxation in a differentiated oligopoly

Diego S. Cardoso*

Department of Agricultural Economics, Purdue University

October 2023

Abstract

Corrective environmental taxes are typically designed to match the value of marginal damages. This approach maximizes social welfare when an environmental externality is the only market imperfection. When multiple imperfections exist, however, knowing the marginal damages is insufficient for setting an optimal tax: it is also necessary to understand the market structure and quantify the effects of imperfections. This paper examines second-best carbon taxation in the US domestic aviation sector and documents novel empirical evidence of the large distortive effects of market imperfections on the efficiency of carbon taxes. Combining sufficient statistics and structural modeling approaches, I estimate abatement costs, calculate the optimal carbon tax under market power and non-environmental taxes, and investigate the effects of introducing a revenue-neutral carbon tax. I find that the current marginal abatement cost is \$211–\$244/ton CO_2. Hence, any positive carbon tax would decrease social welfare if the social cost of carbon (SCC) is smaller than this value. Under a higher SCC of \$300/ton CO_2, the optimal carbon tax is \$107/ton CO_2, thus much lower than the standard Pigouvian tax. Moreover, attempting to improve policy efficiency with a revenue-neutral carbon tax would not yield a double dividend. While welfare gains would follow from reducing the tax deadweight loss and market power, price effects would be dampened by markup adjustments, leading to a net increase in aggregate emissions.

---

*Email: cardosod@purdue.edu. I am grateful to Ivan Rudik, Cathy Kling, Shanjun Li, Todd Gerarden, Julien Daubanes, Ben Leyden, Alberto Salvo, Victor Aguirregabiria, and Cuicui Chen for their comments and suggestions. I also thank participants at the AERE, CREEA, and EEA Conferences, the Northeast Workshop on Energy Policy and Environmental Economics, and seminar participants at Hamilton College, Institute for Policy Integrity/NYU, Michigan State University, Norwegian School of Economics, Syracuse University, University of Connecticut, University of Geneva, ZEW, and various workshops at Cornell University.

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—enjoy 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. When this assumption holds, optimal environmental taxation only requires knowledge about marginal damages. But numerous polluting sectors deviate from perfect markets. For instance, several polluting markets are oligopolies. As Buchanan (1969) demonstrates, optimal taxes differ from marginal environmental costs when firms have market power. Furthermore, most sectors are subject to non-Pigouvian distortionary taxes that lower equilibrium quantities. These taxes decrease related damages and partially substitute a Pigouvian tax. In those cases, an environmental tax based exclusively on marginal damages can even reduce welfare when the assumption of otherwise perfect markets does not hold. Therefore, when multiple imperfections exist, estimates of marginal damage are not sufficient for an efficient policy: it is also necessary to understand the market structure and quantify its imperfections.

This paper combines sufficient statistics and structural modeling approaches to estimate second-best carbon taxes 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 have substantial 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 also 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 global 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 estimates sufficient statistics for welfare changes and structural parameters of a sector model. Based on these estimates, 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.

This paper finds that any positive carbon tax would decrease social welfare if the social cost of carbon[1] (SCC) is equal to its current reference value of \$50 per metric ton of CO_2. Under a higher SCC of \$300/ton CO_2—an upper-bound estimate from Rennert et al. (2022)—the optimal carbon tax is approximately \$107/ton CO_2. Thus, 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 traced to existing market imperfections.

Estimates from the structural model and sufficient statistics indicate that abating one ton of CO_2 by increasing carbon taxes would lead to a loss of \$154–185 in private surplus due to markups and \$53–54 due to the 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 reduction via taxation, the baseline marginal abatement cost is between \$211 and \$244/ton CO_2.

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

[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 result in the so-called double dividend: reduce emissions and further improve welfare by replacing a more distortive tax. I investigate this alternative in counterfactual analyses that implement a revenue-neutral carbon tax to replace the current sales tax.

In line with theory, I find that a revenue-neutral carbon tax of about \$61/ton CO_2 would increase aggregate welfare even under an SCC as high as \$300/ton CO_2. Surprisingly, however, the double dividend fails here because this tax substitution would increase aggregate emissions. Hence, net welfare gains follow from a private surplus increase that exceeds the additional damages. This result is driven by 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 not enough to offset the total emission increase from the less polluting options.

These findings highlight one of the main challenges for climate policy in the sector: the abatement cost of carbon emissions in aviation is high, at least with limited abatement technologies available in the short run. Moreover, with existing distortions, attempts to reform tax instruments to achieve emission targets might backfire.

This paper makes four main contributions. First, it provides robust empirical evidence that adding a carbon tax to a major oligopolistic sector could decrease social welfare for a range of plausible SCC values. Therefore, this analysis contributes to the understanding of the welfare impacts of environmental policy under imperfect competition. 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 advanced theoretical models to characterize the interplay of environmental externalities and market power (Requate, 2006). However, empirical work on this interaction has emerged more recently, typically focusing on a single industry. For instance, empirical analyses to date have documented how existing market power can lead to emission reductions in electricity (Mansur, 2007) and how environmental regulations can exacerbate market power in cement (Ryan, 2012). Similarly, market power and incomplete cost pass-through have been shown to reduce the efficiency of carbon taxes in coal

(Preonas, 2023); incomplete cost pass-through has also been documented in several other carbon-intensive sectors, such as fossil fuels and manufacturing (e.g., Lade & Bushnell, 2019; Ganapati et al., 2020; Muehlegger & Sweeney, 2022). Contributing to this literature, the present paper quantifies efficiency loss of policy instruments under imperfect competition and provides a comprehensive welfare analysis to inform optimal policy design.

Second, this work quantifies the increase in abatement costs due to non-environmental market distortions and provides empirical evidence that a revenue-neutral carbon tax could lead to higher emissions. In doing so, this article contributes to the literature on second-best externality taxation and the interactions between environmental and non-environmental taxes. A significant part of this literature is concerned with the conditions for a double dividend. Most frequently, second-best settings have considered general equilibrium effects 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). Previous studies have also indicated that distortionary taxes increase the cost of environmental policy (Goulder et al., 1999) and that monopoly power can intensify this effect (Fullerton & Metcalf, 2002). However, increased costs can be attenuated by revenue-neutral tax substitution, even when the double dividend fails to materialize (Goulder, 1998). This paper contributes to this literature by providing empirical evidence that emissions can increase in a differentiated oligopoly even when net welfare gains exist. Therefore, these results illustrate a possible conflict between carbon taxes and emission targets, with implications for climate policy choices.

Third, this paper contributes to the literature on empirical welfare analysis by demonstrating how sufficient statistics and structural approaches can be combined to assess marginal and non-marginal effects of environmental policies in imperfect markets. The sufficient statistics approach, introduced by Chetty (2009), evaluates marginal welfare changes from a policy based on a small set of key parameters that can be identified through reduced-form estimation. For instance, previous studies on environmental policies have used sufficient statistics to estimate the efficiency costs of price instruments (Jacobsen, Knittel, Sallee, & Van Benthem, 2020) and characterize market shifts to quantify 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 leverages sufficient statistics to assess the role of non-environmental market distortions and estimate baseline marginal abatement costs. In doing so, this paper also illustrates how these statistics offer a local consistency check for predictions of a structural model necessary for evaluating non-marginal changes.

Fourth, this article offers a novel benefit-cost analysis of prospective climate policy for the US domestic aviation sector. Previous studies have primarily focused on predicting the effects of hypothetical policies on prices and demand. This literature has examined, for example, the effects of a carbon taxes of \$40 (Brueckner & Abreu, 2017) and \$50/ton CO_2 (Pagoni & Psaraki-Kalouptsidi, 2016), cap-and-trade programs (Winchester et al., 2013), and small increases in the jet fuel tax (Fukui & Miyoshi, 2017). However, to date, studies in this literature have overlooked welfare consequences and the role of non-environmental market imperfections.

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 paper's theoretical framework and derives expressions to characterize welfare changes, market distortions, and optimal taxes. A model of the US aviation sector is outlined in Section 4. Section 5 describes the data used in this paper, while section 6 explains the estimation procedures and discusses estimated parameters. The estimation of optimal taxes and the impacts of an aviation carbon tax 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.[2] We can broadly organize airlines 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 be-

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

yond 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; Aryal et al., 2022), entry deterrence (Ciliberto & Williams, 2010; Aguirregabiria & Ho, 2012; Ciliberto & Zhang, 2017), and mergers and consolidation (Kim & Singal, 1993). In the 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 CO_2 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 regulation of carbon emissions, the sector lags behind other industries in decarbonization. As a result, aviation may account for up to 22% of total CO_2 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 Theoretical framework

In this paper, a carbon tax is considered optimal in a second-best sense: it maximizes social welfare, taking market power and distortionary revenue-raising taxes as given. This section introduces a framework for evaluating the welfare effects of a carbon tax and uses that framework to characterize the optimal tax and the role of non-environmental distortions. Throughout the paper, all analyses are made in partial equilibrium, focusing only on the markets that generate the environmental externality of interest.

3.1 An illustration of welfare effects

Before proceeding with a formal analysis, let us graphically illustrate a simplified case. As in Buchanan (1969), consider a single-product monopolist. This single product has a constant marginal cost of production MC, and its consumption generates a marginal externality ϕ. Moreover, the demand function for this good is linear.

Panel (a) in Figure 1 extends the original diagram in Buchanan (1969) by adding a sales tax, which creates a wedge between the inverse demand curve (D) and the price received by the monopolist ($\tilde{P}$). If this market satisfies perfect competition, the competitive equilibrium price (P_c) will equal the marginal private cost (MC). Then, the externality will generate a dead-weight loss (area a), which could be corrected by levying a per-unit tax equal to ϕ; i.e., it would achieve the standard Pigouvian setting at the efficient equilibrium (Q_c^*, P_c^*).

However, when the firm is a monopolist and a sales tax exists, the initial market equilibrium is at Q_{mt}, P_{mt}. Introducing a tax ϕ leads the monopolist to decrease supply even further, to Q'_{mt}, with a higher equilibrium price P'_{mt}. As a result, the interaction

Figure 1: Pigouvian and optimal taxes with multiple market distortions in a polluting monopoly.

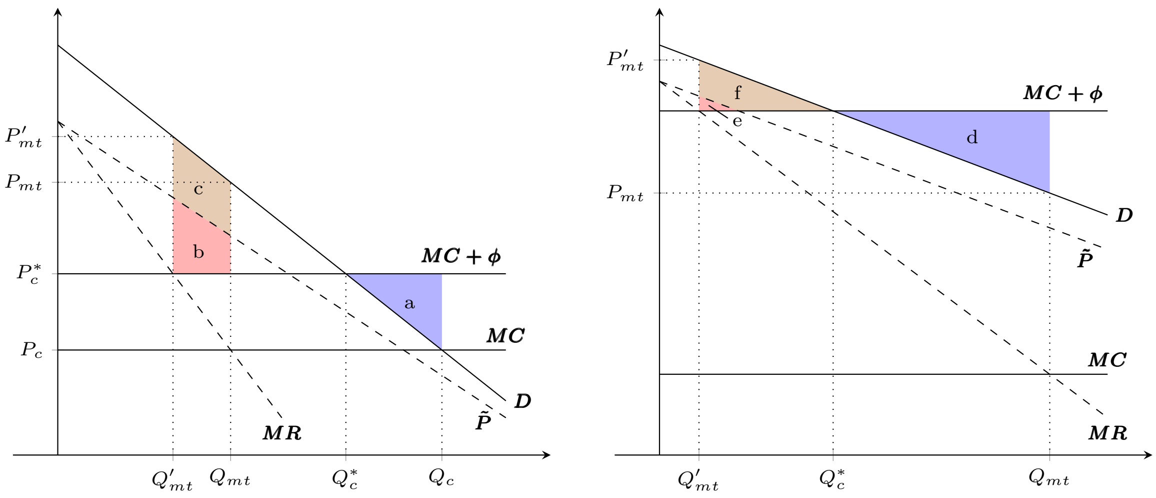

(a) Welfare-decreasing Pigouvian tax

(b) Welfare-improving Pigouvian tax

Notes: D is demand, $\tilde{P}$ is price received (before taxes), MR is marginal revenue, MC is marginal private cost, and ϕ is marginal damage from pollution. Area a is welfare gain from a Pigouvian tax under perfect competition and d under imperfect competition. Areas b and e are welfare losses from monopoly power. Areas c and f are welfare losses from distortionary taxation.

between the externality tax and other market distortions decreases welfare for two reasons. First, market power leads to the reduction represented by area b; this is the welfare loss identified by Buchanan (1969). Second, sales tax distortion drives the loss represented by area c. Hence, for the case illustrated in panel (a) of Figure 1, the standard Pigouvian tax would decrease social welfare. In fact, since $Q_{mt} < Q^*_c$, any positive externality tax would lead to a welfare loss.

The case represented in panel (a) of Figure 1 is a particular one: the externality is small relative to other distortions. Panel (b) illustrates a different scenario, where the Pigouvian tax still improves welfare, but not efficiently. In this case, the Pigouvian tax corrects the externality and increases welfare proportionally to the area d. However, this tax leads the market equilibrium to Q'_{mt}, below the efficient level. As a result, there are welfare losses corresponding to areas e and f (analogous to b and c in Figure 1). The net effect on welfare can be positive as long as the gains from the corrected externality exceed losses from the other distortions.

When the externality is large relative to other distortions, an efficient alternative to the Pigouvian tax is a smaller tax τ^*. This optimal tax leads the monopolist to supply at the efficient market outcome, Q^*_c. Hence, in this particular case, a social planner could efficiently correct the environmental externality under market power and sales tax distortions.

3.2 Social welfare in a taxed oligopoly

The intuitive results obtained from the graphical analysis above can be formalized. Consider a market with $K_m + 1$ differentiated goods indexed by k. Consumption of goods $k = 1, \ldots, K_m$ generates carbon emissions e_k per unit consumed. Index $k = 0$ represents a composite consumption good capturing the "outside option." This composite good has a unitary price and does not generate consumption externalities. Furthermore, no emission abatement technologies exist in the short run, so all abatement is achieved through consumption reduction.

In this setting, social welfare comprises four components: consumer surplus, firms' operating profits, tax revenue,[3] and environmental damages. Throughout the paper, *short-run private surplus* (SRPS) refers to the unweighted sum of consumer surplus, operating profits, and tax revenue.

Consumer surplus. There are N_m identical consumers with quasi-linear utility $U\left(x_0, x_1, \ldots, x_{K_m}\right) = \alpha x_0 + \sum_{k=1}^{K_m} u_k\left(x_k\right)$, where x_k represents the quantities consumed and α determines the marginal utility of income. Under standard assumptions about the utility function, utility maximization implicitly defines demands for each good as $x_k^* = x_k\left(\boldsymbol{P}_m, y\right)$, where $\boldsymbol{P}_m = \left(1, p_1, \ldots, p_k\right)$ is the vector of prices and y is the consumer's income level. Let $q_k\left(\tau\right)$ and $p_k\left(\tau\right)$ represent the equilibrium quantity and price with emissions tax τ. Then, the aggregate money-metric consumer surplus in this market can be represented as

$$CS_m = \frac{N_m}{\alpha} \sum_{k=1}^{K_m} u_k\left(\frac{q_k\left(\tau\right)}{N_m}\right) - \sum_{k=1}^{K_m} p_k\left(\tau\right) q_k\left(\tau\right) + N_m y. \tag{1}$$

Operating profits. There are J_m firms supplying K_m emission-generating goods; the outside option is competitively supplied and can be abstracted from profit considerations. Each product k is subject to a uniform sales tax r and a product-specific

[3] 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.

lump-sum fee ι_k. Thus, $p_k = (1+r)\,\tilde{p}_k + \iota_k$, where $\tilde{p}_k$ is the pre-tax price received by the firm. In this framework, fees ι_k are understood as infrastructure costs paid by the consumer. In aviation, these costs refer to airport and security services provided by entities other than the airlines. These fees are not taxes but must be included in the model because they create a wedge between the ticket price paid and the amount received by airlines, thus affecting pricing behavior. Let $\tilde{p}_k(\tau)$ represent equilibrium pre-tax prices and c_k represent the respective marginal operating costs of production, here assumed constant. Then, total operating profits[4] can be written as

$$\Pi_m(\tau) = \sum_{k=1}^{K_m} \left[\tilde{p}_k(\tau) - c_k - \tau e_k\right] q_k(\tau). \tag{2}$$

Tax revenue. There are two sources of tax revenue: a sales tax r and an emissions tax τ. Total tax revenue, then, is given by

$$T_m(\tau) = \tau \sum_{k=1}^{K_m} q_k(\tau)\, e_k + r \sum_{k=1}^{K_m} q_k(\tau)\, \tilde{p}_k(\tau). \tag{3}$$

Environmental damage. Each unit of emission produces a constant environmental damage ϕ. For carbon emissions, these damages can be understood as the present value of the stream of future damages, i.e., the social cost of carbon. Under this setting, environmental damages in this market are given by

$$\Phi_m(\tau) = \phi \sum_{k=1}^{K_m} q_k(\tau)\, e_k. \tag{4}$$

Combining all four components yields the expression for social welfare in a market:

$$W_m(\tau) = \sum_{k=1}^{K_m} \left\{ \frac{N_m}{\alpha} u_k\left(\frac{q_k(\tau)}{N_m}\right) - \left[\iota_k + c_k + \phi e_k\right] q_k(\tau) \right\} + N_m y. \tag{5}$$

As usual, tax revenues are transfers that cancel out. Thus, social welfare is a function of the utility derived from each good and the private and external costs of providing these goods. Note that the rightmost term in (5) is constant and does not affect the

[4] These are, in fact, variable operating profits. Entry and exit decisions driven by tax changes are not considered in this short-run analysis, for which reason I omit fixed costs.

assessment of welfare changes.

3.3 Marginal welfare effects

Assuming that a marginal change in taxes has negligible effects on costs and markups, differentiating (5) with respect to τ gives the expression for marginal welfare change:

$$\frac{dW_m}{d\tau} = \sum_{k=1}^{K_m} \left[r\tilde{p}_k + \mu_k + (\tau - \phi)\, e_k \right] \frac{dq_k}{d\tau}. \tag{6}$$

The terms within square brackets demonstrate how all three market imperfections affect welfare. In particular, the terms referring to sales tax and markup distortions are analogous to the usual Harberger triangle terms (Kleven, 2021). For these terms, the welfare losses are equal to the "mechanical variation" in tax revenue and operating profit (i.e., holding prices constant). The third term inside the square brackets refers to the environmental externality and its correction mechanism.

The standard Pigouvian taxation arises as a special case in (6) when there is no sales tax nor market power ($r = \mu_k = 0$). The usual prescription applies in this case: setting the environmental tax to marginal damages ($\tau = \phi$) maximizes social welfare (provided that second-order conditions are satisfied). In contrast, if $\tau = \phi$ but other distortions exist, then

$$\left.\frac{dW_m}{d\tau}\right|_{\tau=\phi} = \sum_{k=1}^{K_m} \left[r\tilde{p}_k + \mu_k \right] \frac{dq_k}{d\tau},$$

which is typically negative because $\frac{dq_k}{d\tau} < 0$ for most goods. This demonstrates that, when other distortions exist, the second-best tax is smaller than the standard Pigouvian tax.

3.4 Marginal abatement cost and second-best carbon tax

Combining the marginal change in SRPS and aggregate emissions ($\sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}$), the marginal abatement cost (MAC) is given by

$$MAC\,(\tau) \equiv \tau + \frac{\sum_{k=1}^{K_m} \left[r\tilde{p}_k + \mu_k \right] \frac{dq_k}{d\tau}}{\sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}}. \tag{7}$$

Abatement costs in this framework refer to the private welfare losses following emission reductions induced by environmental tax τ. Equation (6) can be used to characterize the optimal environmental tax with other distortions. The first-order condition for the optimal tax τ^* yields

$$\tau^* = \phi - \underbrace{\left\{ \frac{\sum_{k=1}^{K_m} \mu_k \frac{dq_k}{d\tau}}{\sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}} + \frac{\sum_{k=1}^{K_m} r\tilde{p}_k \frac{dq_k}{d\tau}}{\sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}} \right\}}_{\text{Tax wedges}}, \tag{8}$$

thus making explicit how non-environmental distortions create a wedge between the optimal environmental tax and marginal damages. In particular, this wedge is determined by the marginal changes in each distortion relative to the marginal change in emissions. Throughout the text, I refer to these two components as *tax wedges*, measured in \$/ton CO_2. Combining (7) and (8), a standard result follows: $MAC(\tau^*) = \phi$. That is, the MAC at the optimal emission level is equal to marginal damage—even though the optimal tax rate is smaller than the marginal damage.[5]

3.5 Tax substitution and the double dividend

The existence of a distortive sales tax opens the possibility of a double dividend: substituting the sales tax for an externality tax could both decrease emissions and improve welfare. In this case, it is assumed that the current level of tax revenues cannot be reduced, as it is used to fund necessary government expenditures. This implies a marginal change in the sales tax rate (dr) for every change in the externality tax ($d\tau$) such that the total tax revenue is unchanged: $dT = 0$. Differentiating (3) for all markets yields

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

This expression shows that the marginal rate of substitution between both taxes depends on the ratio of aggregate emissions and revenue, both weighted by the marginally adjusted demands. To examine whether the double dividend materializes, we can as-

[5]This characterization also assumes a uniform tax. In theory, a product-specific tax schedule would weakly improve the outcome of this market. Nevertheless, it is difficult in practice to implement taxes that are specific to each firm or product, especially for input taxes such as the one studied in this paper.

sess marginal changes in welfare (equation 6) and emissions ($dE = \sum_k e_k \frac{dq_k}{d\tau}$). In this scenario, however, raising τ has two effects: the direct effect of a higher emissions tax and the indirect effect of a lower sales tax. Under the assumptions above, it follows that

$$\frac{dq_k}{d\tau} = \frac{\partial q_k}{\partial \tau} + \frac{\partial q_k}{\partial r}\frac{dr}{d\tau} = \left[(1+r)\, e_k + \tilde{p}_k \frac{dr}{d\tau}\right] \frac{dq_k}{dp_k}. \tag{10}$$

Therefore, whether the demand for each product increases or not depends on the relation between emission intensity and pre-tax prices. Plugging (9) into (10), we have

$$\frac{dq_k}{d\tau} \leq 0 \Leftrightarrow \frac{r\tilde{p}_k}{e_k} \leq \frac{\sum_j r\tilde{p}_j \left[q_j + (\tau e_j + r\tilde{p}_j)\, \frac{dq_j}{dp_j}\right]}{\sum_j e_j \left[qj_k + (\tau e_j + r\tilde{p}_j)\, \frac{dq_j}{dp_j}\right]}. \tag{11}$$

Here, $\frac{r\tilde{p}_k}{e_k}$ is the carbon tax implied by the sales tax paid. This expression shows that demand decreases for products with an implied tax above the average implied tax weighted by marginally-adjusted quantities. Since a product's price is not necessarily proportional to its marginal surplus and emissions, we cannot ascertain a priori the direction of the effects of a tax substitution. Therefore, changes in welfare and emissions might go either way, depending on the relative weights of products—whether the double dividend holds is an empirical question.

4 A model of commercial aviation

Marginal welfare analyses can shed light on the effects of small changes in the jet fuel tax. These effects can be approximated using a sufficient statistics approach, as described in section 6.1. To estimate the optimal tax, however, it is necessary to evaluate non-marginal tax changes—this is where the sufficient statistics approach proves limited (Kleven, 2021). Non-marginal changes can be estimated by parameterizing the market equilibrium with structural modeling. In this section, I outline a model for US domestic aviation. This model builds on previous studies on the aviation sector, such as Berry et al. (2006), Berry and Jia (2010), Aguirregabiria and Ho (2012), and Pagoni and Psaraki-Kalouptsidi (2016).

4.1 Definitions

In this model, time is discrete, and each period represents a quarter. A *location* is a city or metropolitan area with one or more airports. A *market*, indexed by m, is a directional pair of locations (of the form $origin \rightarrow destination$), as inBerry et al. (2006) and Aguirregabiria and Ho (2012).[6] A *segment* is an ordered pair between two airports. A *route* r is a sequence of up to four segments forming a round trip. Routes are represented by a four-tuple $(a_o, a_{c1}, a_d, a_{c2})$ of airports: the origin, the outbound connection (if any), the destination, and the inbound connection (if any).[7]

A *product* in this industry is a route r operated by airline i at period t. For simplicity of notation, let $k = (r, i, t)$ index products. Let $\mathcal{K}$ represent the set of available products, using subscripts to indicate partitions. For example, $\mathcal{K}_{mt}$ is the set of products in market m at time t, while $\mathcal{K}_{imt}$ contains only airline i's products in $\mathcal{K}_{mt}$.

4.2 Consumers

There are N_m consumers in market m. At each period, consumers decide whether to purchase at most one of the products available in this market. For consumer n, purchasing product k yields utility as follows

$$u_{nk} = \underbrace{X_k^D \beta^D - \alpha p_k + \xi_k}_{V_k} + \nu_n(\lambda) + \lambda \epsilon_{nk},$$

where X_k^D is a vector of observed characteristics for product k, p_k is the ticket price, and ξ_k represents unobserved (in the data) product characteristics; α, λ, and β^D are model parameters. To simplify notation, $V_k \equiv X_k^D \beta^D - \alpha p_k + \xi_k$ represents the average consumer surplus for product k. The average surplus for the choice of not purchasing a product, indexed by $k = 0$, is normalized to zero.

Consumer-specific tastes are represented by the additive error term $\nu_n(\lambda) + \lambda \epsilon_{nk}$,

[6] Pagoni and Psaraki-Kalouptsidi (2016) have a similar approach, but define locations as a cluster of airports within a radius. In contrast, other studies have defined markets as directional airport pairs (e.g. Borenstein, 1989; Ciliberto & Tamer, 2009; Berry & Jia, 2010). In this paper, markets are defined over metro areas to allow the model to capture competition between flights departing from airports in close proximity.

[7] This model does not consider flights with disjoint segments or with more than one connection each way. In the data used for estimation, these excluded flights correspond to less than 3% of all domestic enplanements.

which yields the nested logit discrete choice model (McFadden, 1978).[8] All flights are grouped in a single nest, denoted by g. The outside option is the single choice available in a separate nest. In this specification, $\nu_n(\lambda)$ is constant across all products and accounts for the correlation of tastes across flights. The term ϵ_{nk} represents the consumer-specific taste for product k. The distribution of the error term is determined by parameter $\lambda \in [0,1]$. When $\lambda = 1$, there is no correlation of random tastes across flights, and the model becomes the standard logit model. When λ approaches 0, the correlation of random tastes goes to 1.

Consumers choose the product that yields the highest utility. Assuming idiosyncratic tastes are distributed Type I Extreme Value, we obtain the probability of a consumer choosing product k. In equilibrium, this probability corresponds to the expected market share of the product, s_k:

$$s_k = Pr\left(u_{nk} \geq u_{nj}, \forall j \in \{0, \mathcal{K}_{mt}\}\right) = \frac{\exp\left(V_k/\lambda\right)}{D_g^{1-\lambda}\left(1 + D_g^{\lambda}\right)}, \tag{12}$$

where $D_g \equiv \sum_j \exp\left(V_j/\lambda\right)$ is the expected utility of purchasing a product in nest g. Under the nested logit specification, the expected consumer surplus in market m at period t is given by

$$CS_{mt} = \frac{1}{\alpha}\ln\left(1 + D_g^{\lambda}\right) + \kappa, \tag{13}$$

where κ is a constant term that is eliminated when evaluating welfare changes (Train, 2009).

4.3 Airlines

In each period and market, airlines maximize operating profits by setting prices for each route they operate in that market. When setting prices, airlines take as given a vector of exogenous demand, cost variables, and the set of routes they operate. Defining the set of routes as given has an important implication for this paper, as it limits

[8] As discussed in section 7, the carbon tax here examined induces rather homogeneous tax changes in the same market because emissions per passenger vary substantially more across than within markets. For this reason, the critical substitution is among nests—all flights vs. not flying—, for which reason this parsimonious demand model adequately captures the main mechanism in a transparent manner.

all analyses to short-run effects only.[9]

Ticket prices paid by consumers include a sales tax (r) and product-specific fees (ι_k). These fees represent the costs of infrastructure services. Taxes and fees create a wedge between the ticket price p_k and the price received by the airline $\tilde{p}_k$ (see section 2 for an overview of these charges).[10] The pre-tax price vector chosen by an airline, $\tilde{P}_{imt} = (\tilde{p}_{k_1}, \tilde{p}_{k_2}, \ldots, \tilde{p}_{k_n})$, maximizes $\Pi_{imt} = \sum_{k \in \mathcal{K}_{imt}} (\tilde{p}_k - c_k) s_k$. The marginal cost per available seat is given by $c_k = \tilde{c}_k + (w_t + \tau) f_k$, where $\tilde{c}_k$ is the constant marginal cost excluding fuel, w_k is the jet fuel cost per gallon, and f_k is the volumetric fuel consumption. The first-order optimality condition for each product k is given by

$$s_k + \sum_{j \in \mathcal{K}_{imt}} (\tilde{p}_j - \tilde{c}_j - (w_j + \tau) f_j) \frac{\partial s_j}{\partial \tilde{p}_k} = 0. \tag{14}$$

The resulting pre-tax price and share vectors, $\tilde{\boldsymbol{P}}_{\boldsymbol{mt}}$ and $\boldsymbol{S}_{\boldsymbol{mt}}$, satisfy a Nash-Bertrand equilibrium. Stacking all first-order conditions from (14), the market equilibrium is a solution to the system of equations

$$\boldsymbol{\mu}_{\boldsymbol{mt}} \equiv \tilde{\boldsymbol{P}}_{\boldsymbol{mt}} - \boldsymbol{C}_{\boldsymbol{mt}} = -J_{mt}^{-1} \boldsymbol{S}_{\boldsymbol{mt}}, \tag{15}$$

where $\boldsymbol{\mu}_{\boldsymbol{mt}}$ is the vector of operating markups, $\boldsymbol{C}_{\boldsymbol{mt}}$ is the vector of marginal operating costs, and J_{mt} is the Jacobian matrix with partial derivatives of quantities with respect to prices, multiplied element-wise by an indicator matrix of product ownership (that is, cell ij is equal to 1 if the products i and j are offered by the same airline and 0 otherwise).[11]

In this paper, the aviation carbon tax is implemented as a volumetric uniform tax

[9] Over a longer horizon, airlines make plans that affect their networks and the markets they serve. Beyond dynamic profit maximization, these decisions take into account strategic considerations and numerous technical and non-technical constraints (Belobaba et al., 2015). Modeling airline decisions to assess long-term network changes requires a substantially more complex model, for which reason long-run effects are left outside the scope of this paper.

[10] Prior literature has largely overlooked the role of taxes and fees in this sector. Nevertheless, most studies were interested in structural characteristics, such as hub premia (Borenstein, 1989, 1991; Ciliberto & Williams, 2010), or large sector shocks, such as mergers (Ciliberto & Tamer, 2009; Aguirregabiria & Ho, 2012), for which these wedges are likely inconsequential. For smaller changes in costs, as the ones considered in the present paper, modeling taxes and fees is essential for better capturing pricing decisions.

[11] Appendix A.3 presents details on how to construct this matrix and solve equilibria given by equation (15).

on jet fuel. Even though isomorphic alternatives exist, such as tradable emission permits and taxes on air travel tickets, I focus on a jet fuel tax for three practical reasons. First, jet fuel is a homogeneous commodity, and jet fuel burn is directly associated with carbon emissions, so a volumetric tax provides a close approximation of the carbon externality. Second, jet fuel is a single-use commodity with limited leakage potential in domestic markets.[12] Third, jet fuel is already taxed in the US; hence, this carbon tax builds on an existing tax structure, which lowers the institutional requirements for its implementation.

5 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, city and metropolitan area populations are collected from the US Bureau of Economic Analysis' Regional Economic Accounts (BEA, 2018).[13]

[12]Though carrying excess fuel from abroad on international flights may be possible, the additional weight increases the fuel burn rate, limiting the economic feasibility of this strategy.

[13]Additional details on the construction of the data set are presented in Appendix C.

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.

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 in this analysis. 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. *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 6. This set of instruments was chosen based on previous studies on the avia-

Table 1: Summary statistics.

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
Price ($)	479.06	157.41	56.97	465.33	2,041.33
Number of stops	1.60	0.64	0	2	2
Market distance (miles)	1,371.16	640.40	67.13	1,256.22	2,724.08
Connection extra distance (miles)	273.36	295.94	0	181.0	2,993
Departures per week	15.85	13.54	0.08	13.00	145.69
Delayed flights (%)	17.28	5.77	0.00	16.88	100.00
Destinations from origin	25.09	19.40	1	18	84
Airlines in market	5.38	1.48	1	5	9
Rivals' products in market	57.34	69.44	0	34	604
Rivals' % of nonstop flights	5.00	9.44	0.00	2.53	100.00
Potential legacy entrants	0.05	0.25	0	0	3
Potential LCC entrants	3.01	1.06	0	3	5
Compl. segment density ($\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
Total 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		

tion 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 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.[14]

6 Empirical approach and estimation

This paper combines two approaches to assess the impact of a carbon tax on aviation. First, using estimated sufficient statistics, I evaluate marginal changes to prices, quantities, and welfare following a small incremental change in the current jet fuel tax. As outlined in Section 3, these marginal changes measure the relative size of market power and tax distortions. Without requiring extensive structural assumptions, sufficient statistics can indicate when an increase in the jet fuel tax is welfare-improving. This first approach, however, has limited use for deriving optimal taxes (Kleven, 2021). To calculate the optimal tax, the second approach relies on estimated structural parameters to characterize non-marginal changes to market equilibria. Specifically, I simulate counterfactual equilibria for various levels of carbon tax to search for the optimal level and analyze its components. In this section, I describe the methods used to estimate parameters and present the estimation results.

6.1 Sufficient statistics

The right-hand side of equation (6) shows two types of terms that are not observed in the data: product-specific markups (μ_k) and marginal quantity changes ($\frac{dq_k}{d\tau}$). These product-specific terms cannot be directly estimated because there is no publicly available data on product-level costs to inform markup calculations. To address this limitation, I focus instead on marginal effects aggregated across products and markets.

For markup data, I use airline average operating revenue and costs per available seat-mile (ASM). These system averages, calculated using Form 41 Financial data, are widely used in the aviation sector to analyze airline performance. Average markups per airline are shown in Table 4. Since these metrics are relative to travel distance, markups vary across flights from the same airline.

Estimating marginal changes in equilibrium outcomes requires further assump-

[14]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 C for details.

tions. First, I assume that a small increase in the fuel tax affects only the fuel price paid by airlines. Therefore, changes to markup, non-fuel costs, and average fuel intensity are negligible following a small fuel tax increment. Under these assumptions, marginal increments in fuel tax have complete pass-through, and the marginal tax shock can be represented as

$$\frac{dp_k}{d\tau} = \frac{dp_k}{dF_k}\frac{dF_k}{d\tau} = \eta_k \frac{p_k}{w_k}, \tag{16}$$

where $F_k = (w_k + \tau) f_k$ is the fuel expenditure per available seat (other variables are defined in section 4.3) and $\eta_k \equiv \frac{\partial \ln p_k}{\partial \ln F_k}$ is the pass-through elasticity. In equilibrium, changes in quantities are functions of the vector of all price changes, and can be expressed as

$$\frac{dq_k}{d\tau} = \sum_{j \in \mathcal{K}_{mt}} \frac{\partial q_k}{\partial p_j}\frac{dp_j}{d\tau} = q_k \sum_{j \in \mathcal{K}_{mt}} \varepsilon_{kj} \frac{\eta_j}{w_j}, \tag{17}$$

where ε_{kj} is the elasticity of demand for product k with respect to the price of product j.

Estimating ε_{kj} for every pair of products in each market is infeasible with the data set at hand. Instead, I estimate how changes in average fuel costs affect the aggregate demand for flights (Q_{mt}). This approach further assumes that small changes in the fuel tax have negligible effects on the relative market shares. Thus, quantities change at the same proportion: $dq_k = s_{k|g} dQ_{mt}$. Flights in a market can then be treated as a composite product $Q_{mt} = \sum_{k \in \mathcal{K}_{mt}} q_k$, with average prices ($p_{mt}$), fuel use ($f_{mt}$), and markups ($\mu_{mt}$) weighted by market shares ($s_{k|g}$). Changes in aggregate quantities can then be expressed as[15]

$$\frac{dQ_{mt}}{d\tau} = \varepsilon \frac{\eta_{mt}}{w_{mt}} Q_{mt}, \tag{18}$$

where and $\varepsilon = \frac{\partial \ln q_{mt}}{\partial \ln p_{mt}}$ is the sector-average elasticity of aggregate demand with respect to the market average ticket price and $\eta_{mt} = \frac{\partial \ln p_{mt}}{\partial \ln \tau}$ is the market-specific, average fuel cost pass-through elasticity.

In this formulation, ε and η_{mt} are the sufficient statistics to be estimated.[16] The suf-

[15]See Appendices A.1 and A.2 for details on the derivation of sufficient statistics and how they are used to calculate marginal welfare changes.

[16]With a focus on public finance, Adachi and Fabinger (2022) derive sufficient statistics for the welfare effects of an *ad valorem* tax under in a general model of imperfect competition. The two sufficient statistics in that setting are tax incidence and the marginal value of public funds, which are functions of

ficient statistic ε is relevant to calculate aggregate marginal changes in prices, quantities, and the components of social welfare. However, it does not affect the estimation of marginal abatement costs or tax wedges. As equation (7) shows, terms that are uniform across markets cancel out. Note that his argument would also apply to any assumed (or estimated) uniform rate of incomplete pass-through.

Estimation and identification. The elasticity of aggregate demand is estimated in a reduced-form approach with the following equation

$$\ln Q_{mt} = \varepsilon \ln p_{mt} + \beta^{(\varepsilon)} X_{mt} + \gamma_m^{(\varepsilon)} + \gamma_t^{(\varepsilon)} + \nu_{mt}^{(\varepsilon)}, \tag{19}$$

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; $\nu_{mt}^{(\varepsilon)}$ is an idiosyncratic demand shock. The fixed effects are added to capture demand components related characteristics of the endpoint locations and seasonality.

As usual in demand estimation, market price p_{mt} is potentially endogenous. To address this source of bias, I construct instruments based on the aviation literature (e.g., Berry et al., 2006; Pagoni & Psaraki-Kalouptsidi, 2016) that capture variations in competition and costs. Competition-shifter instruments include the number of potential legacy entrants and the number of potential low-cost entrants. The cost-shifting instrument is the average fuel expenditure per available seat. The identification mechanism relies on the assumption that consumer decisions are not based on the threat of potential airline competitors or changes in fuel costs, except through the effect of these mechanisms on prices.

Market-specific fuel cost pass-through elasticities (η_{mt}) are calculated directly from the data. Under complete marginal pass-through, this elasticity is equal to the ratio of fuel cost share to ticket price and can be calculated using observed variables: $\eta_{mt} = (1 + r) F_{mt}/p_{mt}$.

Estimation results. Table 2 shows the results of estimating equation (19) 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

pass-through rates and demand elasticities.

Table 2: Results for the estimation of sufficient statistics.

	OLS ln(Passengers)	2SLS, 2nd stage ln(Passengers)	1st stage ln(Price)
ln(Price) [ε]	−0.472	−1.797	
	(0.055)	(0.136)	
Products in market	0.015	0.015	0.0004
	(0.001)	(0.001)	(0.0001)
% of nonstop flights	−0.059	−0.246	−0.054
	(0.146)	(0.205)	(0.061)
Instruments			
log(Fuel expenditure)			0.751
			(0.043)
Potential legacy entrants			0.009
			(0.003)
Potential LCC entrants			−0.012
			(0.007)
Fixed effects			
City pair	Yes	Yes	Yes
Quarter	Yes	Yes	Yes
Observations	19,750	19,750	19,750
R^2	0.963	0.959	0.921
F-statistic			79.2

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

and in line with the literature. For comparison, Berry and Jia (2010)—one of the few papers reporting aggregate price elasticities—find that these elasticities have been increasing in absolute value over time, reflecting structural and demand changes, including quality of service and consumer behavior. Based on the estimated structural parameters of an industry model, they calculate a sector-aggregate elasticity of -1.55 for 1999 and -1.67 for 2006. The demand elasticity here estimate is about -1.80, thus larger in absolute value than those reported in Berry and Jia (2010). Part of the difference in estimates can be attributed to a continuation of the trends identified in that paper, although the difference is relatively small.

6.2 Structural parameters

Demand specification. Manipulating equation (12), we obtain the estimating equation for the nested logit demand (Berry, 1994)

$$\ln s_k - \ln s_0 = X_k^D \beta^D - \alpha p_k + (1-\lambda) \ln s_{k|g} + \xi_k, \tag{20}$$

from which parameters β^D, α, and λ are estimated. Vector X_k^D includes the following observed product characteristics: (i) service frequency in departures per week, (ii) number of stops, (iii) market distance (i.e., between endpoint cities) 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 airline from the origin city, and (vii) fixed effects for airline, origin city-by-quarter, and destination city-by-quarter.

Supply specification. Using estimated parameters $\widehat{\alpha}$ and $\widehat{\lambda}$, observed prices, and observed market shares, we can calculate predicted marginal operating costs by rewriting equation (15) as $\widehat{\boldsymbol{C}}_{\boldsymbol{mt}} = \tilde{\boldsymbol{P}}_{\boldsymbol{mt}} + \widehat{J}_{mt}^{-1} \boldsymbol{S}_{\boldsymbol{mt}}$. These predicted costs are then used to estimate the supply side of the model.

Understanding the cost structure of commercial aviation is critical for the correct specification of the supply-side equation. I present next a brief overview of this structure, as described in Belobaba et al. (2015). Flight operating costs can be mapped into five categories based on their respective unit of variation. First, there are *costs per block hour*. This category includes all aircraft operating costs, which are directly proportional to the time an airplane is used; it also includes passenger service costs, such as flight attendant wages, entertainment, and food, which are proportional to the duration of a flight. Second, there are *costs per departure*, which are primarily aircraft servicing costs; these include cleaning, fueling, and related ground operations. Third, there are *costs per enplaned passenger*, which account for traffic servicing costs, such as passenger and baggage processing. Fourth, there are *costs per distance*, which reflect primarily fuel costs. Fifth and finally, there are *indirect and overhead costs*; this category includes sales, advertising, management, and other categories that are not clearly mapped to any specific units of the flight operation.

The specification of the cost equation builds on the different categories described above:

$$\widehat{c}_k = \rho F_k + \beta_i^S \text{Ramp-to-ramp}_k + \gamma_{i,o} + \gamma_{i,c_1} + \gamma_{i,d} + \gamma_{i,c_2} + \gamma_t + \omega_k, \tag{21}$$

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 of 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 in equation (21) is ρ: it informs how the implied cost varies across the different levels of fuel expenditure. The additional terms in equation (21) map the other cost categories described above. Ramp-to-ramp_k is a proxy for costs per block hour, with parameter β_i^S accommodating differences across airlines. A rich set of fixed effects captures how average costs and costs per enplaned passenger and departure vary for each airline at each airport. A time fixed effect captures average variation in costs across quarters.

Estimation. In the aviation literature, there is often a trade-off between the dimensionality of the characteristics space and the computational performance of the estimation procedure. Even though most papers work with large data sets, it is common to use a small set of proxies and dummy variables instead of a flexible set of fixed effects, especially when applying maximum likelihood or generalized method of moments (GMM) estimators. For example, many papers have used average temperatures and dummies for tourist destinations (e.g., Reiss & Spiller, 1989; Berry et al., 2006; Berry & Jia, 2010) and a dummy variable for whether an airport is slot-controlled or a hub (e.g., Berry & Jia, 2010; Pagoni & Psaraki-Kalouptsidi, 2016). One exception is Aguirregabiria and Ho (2012), which specifies demand and cost equations with several fixed effects; that paper, however, performs separate 2SLS estimations for demand and supply, thus adding the assumption that error terms across both equations are uncorrelated.

Specifications with a high-dimensional characteristic space create additional issues, especially for GMM estimation. Since the joint demand-supply system is nonlinear in parameters α and λ, fixed effects cannot be directly factored out of the equations. Leaving a large number of dummy variables raises computer memory requirements and computation time, both of which can increase exponentially with the number of covariates. Moreover, numerical minimization routines become more computationally challenging when there are many moment conditions (Bennett, Kallus, & Schnabel, 2019).

I address the limitation described above using a method proposed by Conlon and Gortmaker (2020).[17] This method first modifies the estimating equations so that nonlinear terms are absorbed in the left-hand side. Then, fixed effects are factored out using the method of alternating projections (Bauschke, Deutsch, Hundal, & Park, 2003). With these transformations, estimating linear parameters becomes computationally simpler. Furthermore, linear parameters can be expressed as functions of nonlinear parameters. These steps result in a much faster estimation since the numerical optimization routine only searches over a 2-dimensional space, with linear parameters calculated in the inner loop.

Identification. Three variables in equation (20) 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) 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 (20).

The exogeneity of fuel expenditures in (21) relies on three factors. First, jet fuel closely follows crude oil price shifts, as shown in Figure A-4. Hence, since jet fuel accounts for a small fraction of refined oil products, shocks to jet fuel prices largely

[17]A detailed description of the estimation procedure is presented in Appendix A.4.

reflect changes that are exogenous to the aviation sector. Second, fleet and network composition affecting fuel use reflect long-term decisions which are unlikely to respond quickly to recent fuel price shocks. Thus, 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). Moreover, behavioral responses to managerial incentives might only marginally affect fuel efficiency. For instance, Gosnell, List, and Metcalfe (2020) find statistically significant efficiency gains from monitoring and incentivizing pilots directly. Nevertheless, 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.

Given the above, identification in this estimation relies on the exogeneity of the instruments in each equation. These identification assumptions can be arranged in a vector of moment conditions of the form

$$\begin{bmatrix} E\left(Z_k^D \xi_k\right) \\ E\left(Z_k^S \omega_k\right) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \tag{22}$$

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—corresponding to approximately 3,000 dummy variables—these conditions form a system with 22 moment conditions and 16 parameters. This system is the base of the 2-step GMM estimation used in this paper.[18] Reported standard errors are heteroskedasticity robust and clustered by non-directional city pairs to allow for correlation across markets in opposite directions.

Estimation results. The results for the joint estimation procedure are shown in Table 3, with demand and supply coefficients in the left and right columns, respectively. The key demand coefficients estimated in Table 3 are largely in line with the literature. Estimates for α in more recent studies generally vary between -1.36 (Aguirregabiria

[18] Appendix B.1 shows the results based on alternative estimators and nesting choices.

Table 3: Results for the joint demand-supply estimation of the structural model.

Demand	$ln(s_{kt}/s_{0t})$	Supply	c_{kt}
Price ($100) [$-\boldsymbol{\alpha}$]	-0.846	Fuel expenses [$\boldsymbol{\rho}$]	0.693
	(0.114)		(0.123)
ln(share within nest) [$\mathbf{1} - \boldsymbol{\lambda}$]	0.400	Total ramp-to-ramp time (h)	0.152
	(0.049)		(0.014)
Departures per week	0.036	× *American*	-0.040
	(0.002)		(0.007)
Number of stops	-0.803	× *Delta*	0.046
	(0.071)		(0.008)
Market distance (100mi.)	0.083	× *United*	-0.018
	(0.015)		(0.008)
Market distance squared (100mi.)^2	-0.001	× *Alaska*	-0.011
	(0.028)		(0.032)
Connection extra distance (100mi.)	-0.084	× *JetBlue*	-0.015
	(0.009)		(0.010)
Connection extra distance sq. (100mi.)^2	0.327	× *Other low-cost*	-0.126
	(0.054)		(0.007)
Share of delayed departure (%)	-0.006		
	(0.001)		
Destinations from origin	0.012		
	(0.002)		
Observations	267,967		
Objective function minimum	3.46 $\times 10^{-5}$		

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.

& 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. Other demand coefficients in Table 3 also present the expected signs and are in line with the literature. For instance, consumers value a higher frequency of service, meaning more opportunities for convenient travel times. 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 tend to avoid flights that were more frequently delayed in the previous quarter, although the effect is rather small.

The supply side of Table 3 shows that a \$1 increase in fuel costs per available seat translates, on average, to a \$0.69 increase 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 cost parameters indicate an expected pattern: the costs per block hour and passenger among legacy carriers are generally greater than those among low-cost carriers.

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 Financial database. This database contains quarterly aggregate indicators reported by airlines, including revenue and operating costs by available seat mile (ASM). These data are not used to estimate structural parameters, so they are good candidates for performing 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. These metrics are widely used for performance evaluation in the aviation sector and capture the focal point of the model's application: markups.

Table 4 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. 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

Table 4: 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	9.64	9.44	3.84	4.26	27.88
American	14.57	14.65	11.45	10.37	3.12	4.28	19.26
Delta	15.91	15.80	12.57	10.89	3.33	4.92	18.93
United	13.91	12.60	11.21	10.75	2.70	1.85	14.30
Other LCCs	4.21	8.31	1.36	6.69	2.86	1.62	8.39
JetBlue	10.73	12.07	8.01	9.46	2.72	2.62	5.62
Alaska	10.07	11.48	7.22	8.37	2.85	3.11	5.61
Sector average	13.03	13.04	9.83	9.58	3.20	3.46	

Notes: Predicted averages result from the estimated sector model. Reported values are calculated based on the BTS 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.

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 least accurate in predicting their revenues.

Errors in revenue prediction, however, are not fully passed onto markup prediction errors 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 though these equilibria may not capture all components of pricing decisions, they reproduce the most important patterns in the reported data and provide an excellent approximation to sector average values.

7 Second-best carbon taxation

Based on the estimated sufficient statistics and structural parameters, this section performs counterfactual analyses to quantify the 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 using both sufficient statistics and structural model approaches. 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/ton CO_2, a reference value commonly used in the literature and in line with the U.S. Federal SCC of \$51 set in 2021. The medium SCC of \$200/ton CO_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/ton CO_2, rounding the estimate of \$308 in Rennert et al. (2022) when a 1.5% discount rate is used.

Each gallon of jet fuel burned emits an average of 9.57 kg CO_2 (EIA, 2016). To account for other greenhouse gases, these emissions can be converted to CO_2-equivalent terms. In this paper, I use a 1.4 conversion factor, which is the central estimate in Azar and Johansson (2012). Hence, one gallon of jet fuel accounts for approximately 13.4 kg of CO_2-equivalent—this is the emission intensity used to calculate climate damages in the welfare analyses below.

7.1 Marginal abatement costs and welfare consequences

Table 5 shows the estimated effects of a 1-cent increment in the current volumetric jet fuel tax based on two different methods. The second column reports marginal effects calculated using sufficient statistics estimated with aggregate market data, as outlined in section 6.1. The third column reports equivalent outcomes calculated by solving the estimated sector model with a 1-cent/gallon tax increase. For reference, this increase is equivalent to an additional carbon tax of about \$0.75/ton CO_2.

Despite the substantial differences in the data used and the assumptions underlying each approach, Table 5 shows that both methods deliver fairly similar estimates for MAC and wedges. However, the simplified models of consumption and product competition underlying sufficient statistics lead to two key differences. First, the lack of substitution within markets underestimates emission reductions relative to the structural approach. This underestimation inflates marginal cost estimates, as it shrinks the denominator. Second, as indicated in Table 4, markups calculated from financial reports are typically higher than those predicted with the structural model, further increasing the gap between MAC estimates due to higher markup wedges. Nevertheless, key estimates remain quantitatively similar: the MAC with sufficient statistics is approximately \$244/ton CO_2, about 16% higher than that estimated with the struc-

Table 5: Effects of raising the jet fuel tax by 1 cent/gallon

	Sufficient statistics	Structural model
Marginal costs (\$/tCO_2)		
Marginal abatement cost (MAC)	244.27	210.92
Markup wedge	185.12	154.27
Sales tax wedge	54.10	53.37
Aggregate market changes		
Mean ticket price	+0.087%	+0.065%
Passengers	-0.156%	-0.197%
Emissions	-0.171%	-0.241%
Welfare changes (million \$)		
Short-run private surplus (SRPS)	-23.8	-29.3
Social welfare, SCC of: \$50	-18.9	-22.3
\$200	-4.2	-1.5
\$300	+5.7	+12.3

Notes: The baseline tax is 4.4 cents/gallon or, equivalently, \$3.28/ton CO_2.

tural model.

The bottom rows in Table 5 show the marginal effects on welfare and allow us to gauge the second-best efficiency of a carbon tax. The private deadweight loss of marginally increasing the tax amounts to about \$24–29 million. Whether private losses are offset by a social benefit of carbon damage reductions depends on the SCC scenario. In the low and medium SCC scenarios, both approaches estimate that private costs exceed the benefits, thus predicting that raising the carbon tax would lead to a net reduction in welfare. In the high SCC scenario, both approaches show welfare gains, indicating that the current jet fuel tax is below its second-best optimal value.

To find the optimal tax level, however, we cannot rely on the sufficient statistics approach, as it evaluates a constant MAC based on marginal welfare changes. In such cases, structural models are better equipped to handle non-marginal assessments (Kleven, 2021). Still, Table 5 illustrates how the comparison of results from these methods serves as a sanity cross-check, as each approach introduces limitations based on their set of assumptions.

7.2 Non-marginal welfare changes and the optimal carbon tax

With a structural model of the sector, it is possible to evaluate non-marginal increases in the jet fuel tax and recalculate the equilibrium outcomes in each market. Panel (a) in Figure 2 shows that social welfare decreases under a low or medium SCC. Hence, no positive optimal tax exists when the SCC is \$50 or \$200. In a high SCC scenario, the optimal tax is approximately \$107/ton CO_2, corresponding to a jet fuel tax increase of \$1.39/gallon. At about a third of the marginal damage, this optimal tax is much lower than the standard Pigouvian tax prescription of \$300.

To motivate the mechanisms driving the optimal tax level, panel (c) in Figure 2 decomposes welfare variations into private surplus and external damages. The marginal damage avoided is analogous to the social benefit of a carbon tax, whereas the marginal loss of SRPS is analogous to the MAC (in terms of private welfare). Note, however, that these marginal values are presented in terms of tax level changes rather than emissions. This panel shows that other market distortions lead to a high marginal loss of SRPS. If damages are low relative to SRPS losses, as in the case of low and medium SCCs, social welfare decreases as the tax increases. When marginal damage is higher than the marginal SRPS loss at the baseline, the optimal tax is found at the intersection of these curves—the standard marginal cost equals marginal benefit result.

Changes in SRPS and emissions allow us to evaluate the costs of abating emissions with a carbon tax. Panel (b) in Figure 2 shows how costs and tax wedges vary with abatement levels. From the baseline, the MAC is approximately \$211/ton CO_2. Hence, with low and medium SCCs, the private welfare cost of abatement exceeds the avoided damages at the margin. When damages are high and exceed the initial MAC, the optimal abatement level is found wherethe MAC curve intersects with the specified SCC level. For instance, the optimal tax for an SCC of \$300/ton CO_2 corresponds to an abatement of 16.6 million tons of CO_2 (a 29% reduction).

Panel (b) in Figure 2 also shows that the wedge due to markups initially accounts for about 75% of the distortion wedges. As the abatement level increases with a higher carbon tax, incomplete pass-through tends to decrease markups, thus reducing the markup wedge. In contrast, with more expensive tickets due to a higher carbon tax, the distortion from the sales tax increases.[19] For the high-damages case, the difference

[19]The estimated MACs consider aggregate abatement costs. However, as market power and car-

Figure 2: Welfare consequences of different levels of carbon taxation.

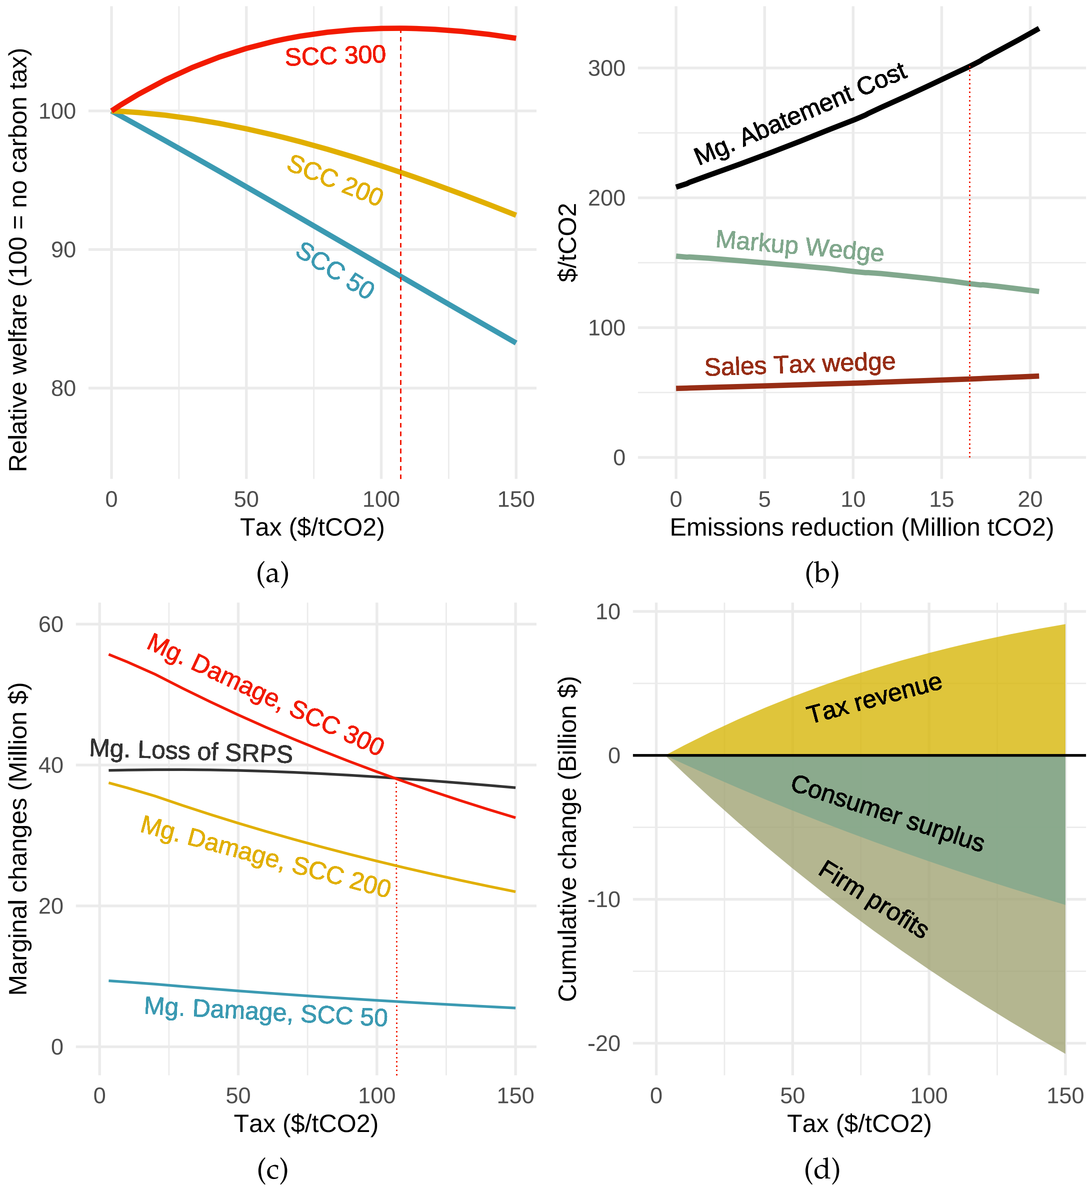

Notes: *Mg. Damage* corresponds to the marginal benefit of lowering emissions, which depends on the social cost of carbon (SCC) and the emission reductions at each tax level. *Loss of SRPS* (short-run private surplus) is analogous to the aggregate private cost of reducing emissions with a carbon tax.

between the $300 SCC and the $107 optimal tax indicates a $193 tax wedge, of which approximately $133 is due to market power and $60 to sales tax.

Abatement costs represent changes in private welfare, which can be further decomposed. Panel (d) in Figure 2 displays the cumulative changes in tax revenue, consumer surplus, and operating profits. This panel shows that the burden of the carbon tax is

bon intensity vary across markets, we can conceptualize market-specific MACs resulting from different wedges. Appendix E examines individual MACs and shows that even the lowest abatement costs are above $100/ton CO_2.

split approximately evenly between firms and consumers. For the high-damage optimal tax of \$107/ton CO_2, consumer surplus would fall by \$3.79B and operating profits by \$3.86B. An increase of \$3.62B in tax revenues, however, could be allocated to partially offset losses to either side.

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. After all, 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 two dividends. First, it could decrease carbon emissions. Second, it can promote net welfare gains by replacing a less efficient taxation scheme. However, as established in section 3.5, we cannot determine a priori whether both dividends exist, 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 \$0.823/gallon corresponds to a carbon tax of \$61.42/ton CO_2.

Table 6 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 unintended consequences for climate concerns. Despite the mean ticket price staying relatively unaffected, a revenue-neutral carbon tax would increase passenger traffic by 1.4% and emissions by 0.7%. The bottom part of Table 6 confirms a revenue-neutral carbon tax increases welfare in all three scenarios. However, these

Table 6: Effects of substituting the existing sales tax with a revenue-neutral carbon tax.

	Changes with a revenue-neutral carbon tax		
Aggregate market changes			
Mean ticket price	+0.04%		
Passengers [million]	+1.43% [+4.02]		
CO_2 emissions [thousand tons]	+0.72% [+850]		
Welfare changes (million $)	SCC of $50	SCC of $200	SCC of $300
Short-run private surplus (SRPS)	+557	+557	+557
Damages	+43	+170	+255
Social welfare	+514	+387	+302

increases are driven by increments in private surplus that exceed the additional damages from slightly higher emissions—i.e., the double dividend does not hold.

At first glance, the results in Table 6 run counter to the intuition that a tax instrument proportional to the externality should reduce damages. The crucial mechanism to understand how a carbon tax could increase emissions lies in the distorting roles of pre-existing sales tax and market power. Figure 3 illustrates these distortions and shows the effects on flights within each quintile of carbon intensity, measured as the emissions per passenger (e_k).

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 3.5.[20] 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. A revenue-neutral carbon tax enforces a uniform tax of 61.42/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 burden of each product. Panel (b) in Figure 3 illustrates this mechanism. It shows that the distribution of changes in taxes collected reflects the adjustment for the initial under/over taxation of emissions. However, the shifts in tax burden are not fully passed through onto prices because

[20] Appendix D 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.

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

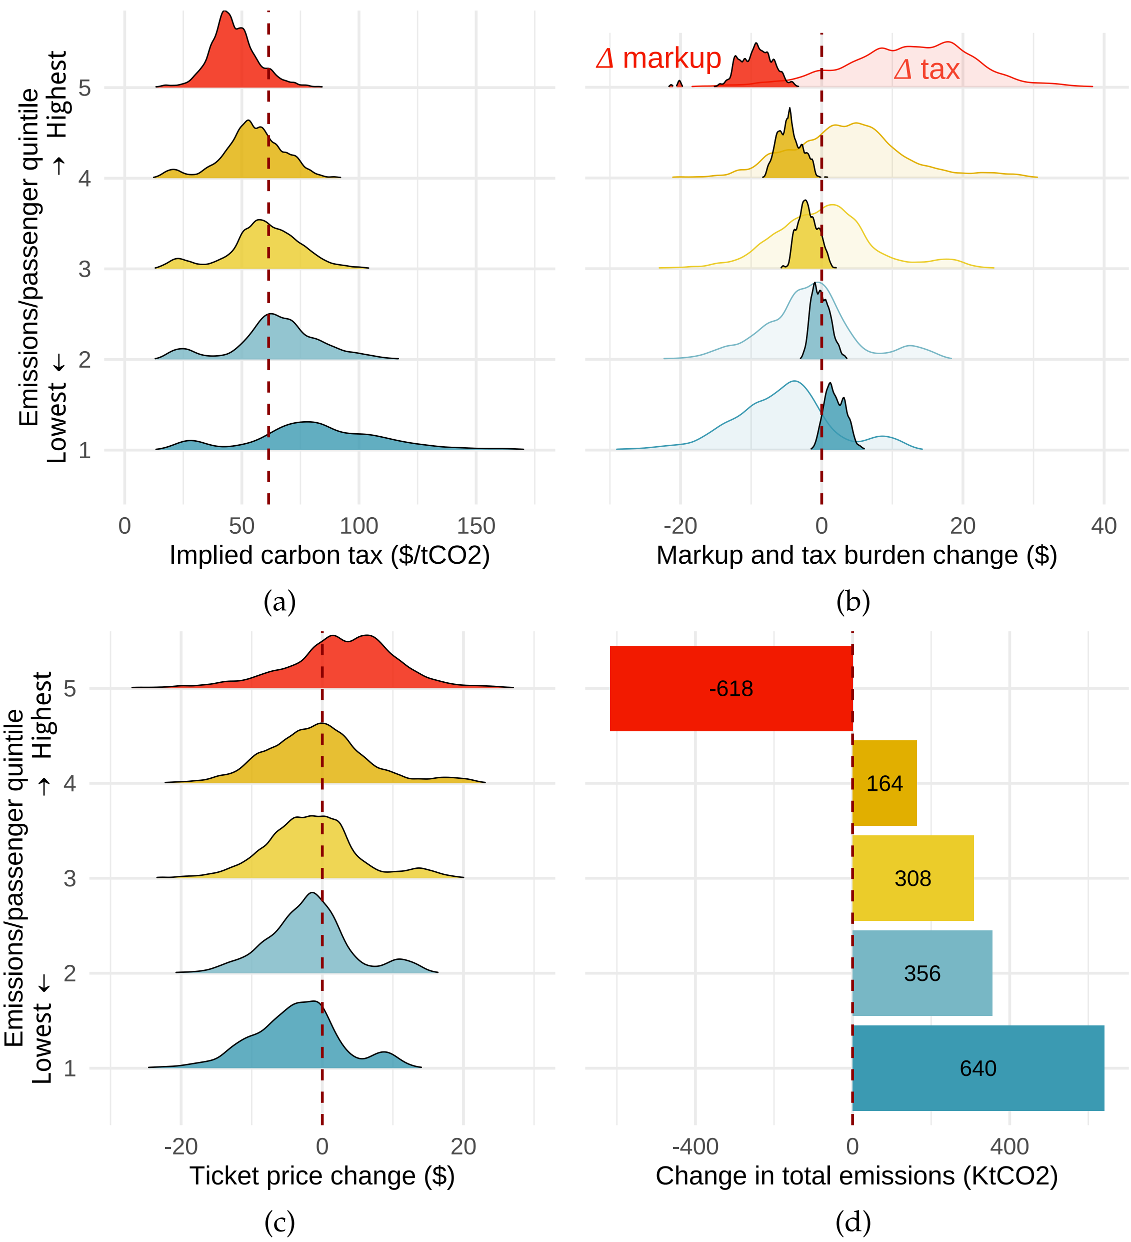

Notes: Curves show kernel densities for each carbon intensity quintile. Distributions are weighted by total emissions, so each quintile has approximately the same baseline aggregate emissions. Quintile cut-offs are approximately 320, 410, 530, and 710 Kg of CO_2 per passenger.

firms adjust markups to compensate for these changes.

Panel (c) in Figure 3 shows that, due to markup adjustments in response to changes in tax burden, a substantial share of emissions in the upper quintiles observes price decreases. The combination of lower markups and the elimination of the deadweight loss of the sales tax contribute to increases in private surplus, as reported in Table 6. However, markup reductions in the upper quintiles effectively shift the price change distribution to the left, leading many carbon-intensive products to experience small reductions in prices. Consequently, the 3rd and 4th quintiles observe net increases in aggregate emissions; the sharp decrease in emissions in the 5th quintile is not enough to offset increases in all other quintiles. Hence, despite net welfare gains, aggregate emissions increase with the tax substitution.

8 Conclusion

This paper has studied how oligopoly market power and existing distortionary taxes affect environmental policy. Building on seminal work in environmental economics (Buchanan, 1969), this paper has shown how market imperfections affect optimal environmental taxes, and how welfare effects can be decomposed and attributed to each market imperfection. Based on this theoretical framework, I have evaluated the impact of a carbon tax on aviation. In doing so, I have used sufficient statistics to calculate marginal effects and a structural estimation approach to calculate non-marginal effects and optimal taxes.

The main findings indicate that existing distortions are large and, at the margin, exceed the climate damages from aviation in scenarios where the social cost of carbon is below $211. As a consequence, there would be no positive optimal carbon tax in this sector unless the social cost of carbon is high. Even if a positive optimal tax exists, it represents a fraction of the marginal damage and, thus, it is below the standard Pigouvian tax prescription. I find that the wedge between marginal damage and optimal tax is primarily driven by market power, which accounts for about three-quarters of the wedge. Further analysis shows that attempting to remediate the distortion brought by an imperfect tax instrument with a revenue-neutral carbon tax could result in increased emissions. This seemingly puzzling result stems from the fact that firms can reduce markups and undo part of the effects of increased taxation on the most polluting flights. In doing so, demand and emissions do not fall enough to decrease aggregate emissions.

The analyses presented in this paper are subject to a number of limitations. First, all results concern short-run effects, as airlines' networks and fleets are held fixed. For this reason, the size of distortions can be underestimated when considering longer periods because higher fuel taxes may lead to exits that would increase market power for the remaining players. However, such decisions typically involve strategic considerations beyond pure short-run profitability (Belobaba et al., 2015). In addition, over longer periods, average fuel efficiency tends to increase slowly due to fleet turnover. Second, all analyses are in partial equilibrium and overlook effects on other transportation modes. In practice, more expensive flights may lead to substitution to other modes on some or all parts of a trip. These substitution possibilities create leakage opportunities, which the models used in this paper do not capture. Third, climate change is the only environmental externality considered in this paper. Aviation has other significant environmental consequences, especially those with local impacts. For instance, fuel burn affects local air pollution and health outcomes (Schlenker & Walker, 2016), and noise pollution affects property values in the neighborhood of airports (Nelson, 2004).

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. These features suggest that alternative policies may be more adequate in the short run. For instance, multi-sector emission permits and offsets could be better alternatives, as they take advantage of lower abatement costs in other sectors. 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 electric planes are potential candidates in development.

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.

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.

Bennett, A., Kallus, N., & Schnabel, T. (2019). Deep generalized method of moments for instrumental variable analysis. In *Advances in neural information processing systems* (pp. 3564–3574).

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.

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., & 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.
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, U. E. P. A. (2022). *Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances.* US Environmental Protection Agency Washington, DC. 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* (Tech. Rep.). Energy Policy Research Group, University of Cambridge. Retrieved 2023-09-22, 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. (1998). Environmental policy making in a second-best setting. *Journal of Applied Economics*, *1*(2), 279–328.

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.

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.

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.

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).

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.

Preonas, L. (2023). Market power in coal shipping and implications for us climate policy. *Review of Economic Studies*, rdad090.

Reiss, P. C., & Spiller, P. T. (1989). Competition and entry in small airline markets. *The Journal of Law and Economics*, *32*(2, Part 2), S179–S202.

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.

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

Contents

A Estimation procedures and welfare analyses

The first two sections in this appendix outline the derivation of sufficient statistics and the corresponding expressions for evaluating marginal welfare changes. The third section shows how market equilibria are solved to calculate implied operating costs and counterfactual outcomes with different tax levels. Finally, the fourth section explains how I estimate the structural model using GMM with high-dimensional fixed effects.

A.1 Derivation of sufficient statistic ε

Consider a composite flight $Q_{mt} = \sum_{k \in \mathcal{K}_{mt}} q_k$ for market m in period t. Let the attributes of the composite flight be the average of attributes among flights in this market, weighted by the relative market shares $s_{k|g}$ (or shares within nest). That is, we define the average pre-tax price $\tilde{p}_{mt} = \sum_{k \in \mathcal{K}_{mt}} s_{k|g} \tilde{p}_k$, average price $p_{mt} = \sum_{k \in \mathcal{K}_{mt}} s_{k|g} p_k$, average fuel use $f_{mt} = \sum_{k \in \mathcal{K}_{mt}} s_{k|g} f_k$, and average markup $\mu_{mt} = \sum_{k \in \mathcal{K}_{mt}} s_{k|g} \mu_k$.

As discussed in section 6.1, I assume that changes to markups, non-fuel costs, and relative market shares are negligible following a small change in tax fuel τ. Therefore, a marginal increment in the tax fuel is fully passed on to consumers.

Summing equation (17) over all products in the market, it follows that

$$\frac{dQ_{mt}}{d\tau} = \sum_{k \in \mathcal{K}_{mt}} \sum_{j \in \mathcal{K}_{mt}} \frac{\partial q_k}{\partial p_j} \frac{dp_j}{d\tilde{p}_j} \frac{d\tilde{p}_j}{d\tau} = (1+r) \sum_{k \in \mathcal{K}_{mt}} \sum_{j \in \mathcal{K}_{mt}} \frac{\partial q_k}{\partial p_j} f_j = (1+r) \sum_{j \in \mathcal{K}_{mt}} f_j \frac{\partial Q_{mt}}{\partial p_j}. \tag{A-1}$$

With no change in relative market shares after a marginal change in τ, it follows that price changes are proportional: $dp_j = s_{j|g} dp_{mt}$. Hence,

$$\frac{\partial Q_{mt}}{\partial p_j} = \frac{\partial Q_{mt}}{\partial p_{mt}} \frac{\partial p_{mt}}{\partial p_j} = s_{j|g} \frac{\partial Q_{mt}}{\partial p_{mt}}. \tag{A-2}$$

Combining (A-1) and (A-2) yields

$$\frac{dQ_{mt}}{d\tau} = (1+r) \sum_{j \in \mathcal{K}_{mt}} s_{j|g} f_j \frac{\partial Q_{mt}}{\partial p_{mt}} = (1+r) f_{mt} \frac{\partial Q_{mt}}{\partial p_{mt}} \tag{A-3}$$

Then, expressing (A-3) in terms of elasticities, it follows that

$$\frac{dQ_{mt}}{d\tau} = (1+r)\frac{f_{mt}}{p_{mt}}Q_{mt}\varepsilon = (1+r)\frac{1}{w_t}\frac{F_{mt}}{\tilde{p}_{mt}}Q_{mt}\varepsilon = \varepsilon\frac{\eta_{mt}}{w_{mt}}Q_{mt}, \tag{A-4}$$

where $\eta_{mt} = \frac{\partial \ln p_{mt}}{\partial \ln w_{mt}}$ is the jet fuel cost pass-through elasticity and $\varepsilon = \frac{\partial \ln Q_{mt}}{\partial \ln p_{mt}}$ is the elasticity of aggregate demand with respect to the market average ticket price.

A.2 Marginal welfare calculations with sufficient statistics

We can use equation (5) to analyze how welfare and its components change given a marginal change in tax τ. Here, τ is a volumetric tax on jet fuel rather than a tax on emissions. Dividing τ by the equivalent carbon intensity of jet fuel burn ($h = 0.0134$ ton CO_2/gallon) gives the tax per ton of CO_2. In this derivation, I assume a new equilibrium exists but remain agnostic about the specific changes in price ($\frac{dp_k}{d\tau}$) and quantity ($\frac{dq_k}{d\tau}$). Differentiating each component with respect to τ (and omitting function arguments for clarity) yields

$$\frac{dCS_m}{d\tau} = -\sum_{k=1}^{K_m} q_k \frac{dp_k}{d\tau} \tag{A-5}$$

$$\frac{d\Pi_m}{d\tau} = \sum_{k=1}^{K_m} \left\{ \mu_k \frac{dq_k}{d\tau} + \left[\frac{d\tilde{p}_k}{d\tau} - f_k \right] q_k \right\} \tag{A-6}$$

$$\frac{dT_m}{d\tau} = \sum_{k=1}^{K_m} \left[q_k f_k + \tau f_k \frac{dq_k}{d\tau} \right] + r \sum_{k=1}^{K_m} \left[q_k \frac{d\tilde{p}_k}{d\tau} + \tilde{p}_k \frac{dq_k}{d\tau} \right] \tag{A-7}$$

$$\frac{d\Phi}{d\tau} = \phi h \sum_{k=1}^{K_m} f_k \frac{dq_k}{d\tau}. \tag{A-8}$$

Equation (A-5) uses the first order optimality condition $u_k'(x_k^*) = \alpha p_k$, for $k \in \{1 \ldots K_m\}$. In (A-6), the term $\mu_k \equiv \tilde{p}_k - \tilde{c}_k - (w_k + \tau) f_k$ denotes the operating markup. In the same equation, $\frac{d\tilde{p}_k}{d\tau} - f_k$ captures the marginal change in markup; this term is equal to zero when there is complete tax pass-through. In (A-8), ϕh is the marginal damage per gallon of jet fuel burned.

We can rewrite equation (A-5) to evaluate a marginal change in consumer surplus using the composite flight Q_{mt}

$$\frac{dCS_{mt}}{d\tau} = -\sum_{k=1}^{K_m} q_k \frac{dp_k}{d\tau} = -(1+r)\, Q_{mt} \sum_{k=1}^{K_m} s_{k|g} f_k = \frac{dCS_{mt}}{d\tau} = -(1+r)\, Q_{mt} f_{mt}. \quad \text{(A-9)}$$

With constant relative market shares under a marginal tax change, it follows that changes in quantities are proportional: $dq_k = s_{k|g} dQ_{mt}$. Then, rewrite equation (A-6) for marginal changes in profits using the composite flight as follows

$$\frac{d\Pi_{mt}}{d\tau} = \sum_{k=1}^{K_m} \left\{ \mu_k \frac{dq_k}{d\tau} + \left[\frac{d\tilde{p}_k}{d\tau} - f_k \right] q_k \right\} = \sum_{k=1}^{K_m} \mu_k s_{k|g} \frac{dQ_{mt}}{d\tau} = \mu_{mt} \frac{dQ_{mt}}{d\tau},$$

where the second step uses the full pass-through assumption to cancel out the term for changes in markups. Then, using (A-4), it follows that

$$\frac{d\Pi_{mt}}{d\tau} = \mu_{mt} \frac{\eta_{mt}}{w_{mt}} Q_{mt} \varepsilon. \quad \text{(A-10)}$$

Following a similar procedure for changes in tax revenue, rewrite equation (A-7) as

$$\begin{aligned} \frac{dT_{mt}}{d\tau} &= \sum_{k=1}^{K_m} \left[q_k f_k + \tau f_k \frac{dq_k}{d\tau} \right] + r \sum_{k=1}^{K_m} \left[q_k f_k + \tilde{p}_k \frac{dq_k}{d\tau} \right] \\ &= (1+r)\, Q_{mt} f_{mt} + \sum_{k=1}^{K_m} (\tau f_k + r\tilde{p}_k)\, s_{k|g} \frac{dQ_{mt}}{d\tau} \\ &= (1+r)\, Q_{mt} f_{mt} + (\tau f_{mt} + r\tilde{p}_{mt}) \frac{dQ_{mt}}{d\tau}, \end{aligned}$$

so that

$$\frac{dT_{mt}}{d\tau} = Q_{mt} \left\{ (1+r)\, f_{mt} + (\tau f_{mt} + r\tilde{p}_{mt}) \frac{\eta_{mt}}{w_{mt}} \varepsilon \right\}. \quad \text{(A-11)}$$

Lastly, for damages, rewrite equation (A-8) as

$$\frac{d\Phi_{mt}}{d\tau} = \phi h \sum_{k=1}^{K_m} f_k \frac{dq_k}{d\tau} = \phi h f_{mt} \frac{dQ_{mt}}{d\tau} = \phi h f_{mt} \frac{\eta_{mt}}{w_{mt}} Q_{mt} \varepsilon. \quad \text{(A-12)}$$

Combining equations (A-9)–(A-12), the marginal change in welfare for market m in period t, using sufficient statistics μ_{mt}, η_{mt}, and ε, is given by

$$\frac{dW_{mt}}{d\tau} = \left(\mu_{mt} + r\tilde{p}_{mt} + (\tau - \phi h) f_{mt}\right) \left[\frac{\eta_{mt}}{w_{mt}} Q_{mt} \varepsilon\right]. \tag{A-13}$$

The change in aggregate welfare for the sector is the sum of (A-13) evaluated for each market and period,

$$\frac{dW}{d\tau} \equiv \sum_{m \in \mathcal{M}} \sum_{t=1}^{4} \frac{dW_{mt}}{d\tau}, \tag{A-14}$$

where $\mathcal{M}$ represents the set of all markets. The marginal change to short-run private surplus is a similar expression, only dropping the damage term:

$$\frac{dSRPS_m}{d\tau} = \left(\mu_{mt} + r\tilde{p}_{mt} + \tau f_{mt}\right) \left[\frac{\eta_{mt}}{w_{mt}} Q_{mt} \varepsilon\right]. \tag{A-15}$$

The marginal abatement cost (MAC) then divides (A-15) by the marginal change in emissions $\frac{de_{mt}}{d\tau} = hf_{mt} \frac{dQ_{mt}}{d\tau}$. Hence,

$$MAC_{mt}(\tau_0) \equiv \frac{\frac{dSRPS_{mt}}{d\tau}}{\frac{de_{mt}}{d\tau}} = \frac{\left(\mu_{mt} + r\tilde{p}_{mt} + \tau f_{mt}\right) \left[\frac{\eta_{mt}}{w_{mt}} Q_{mt} \varepsilon\right]}{hf_{mt} \left[\frac{\eta_{mt}}{w_{mt}} Q_{mt} \varepsilon\right]},$$

so

$$MAC_{mt}(\tau_0) = \frac{1}{h} \left[\tau + \frac{\mu_{mt} + r\tilde{p}_{mt}}{f_{mt}}\right], \tag{A-16}$$

where $\tau_0 = \$0.044$/gallon is the baseline jet fuel tax. Aggregating for the entire sector, the marginal abatement cost is given by

$$MAC(\tau_0) \equiv \frac{\frac{dSRPS}{d\tau}}{\frac{de}{d\tau}} = \frac{1}{h} \left\{\tau + \frac{\sum_{m \in \mathcal{M}} \sum_{t=1}^{4} \left(\mu_{mt} + r\tilde{p}_{mt}\right) \left[\frac{\eta_{mt}}{w_{mt}}\right]}{\sum_{m \in \mathcal{M}} \sum_{t=1}^{4} f_{mt} \left[\frac{\eta_{mt}}{w_{mt}} Q_{mt}\right]}\right\}. \tag{A-17}$$

Note that the expressions for the MAC do not depend on elasticity ε; this reflects the assumption that changes in quantities are proportional, so emissions change linearly with quantities.

A.3 Solving market equilibria and operating costs

As described in section 4.3, firms choose a vector of pre-tax prices $\tilde{\boldsymbol{P}}_{\boldsymbol{imt}}$ that maximize market profits, taking the competitor prices as given

$$\tilde{\boldsymbol{P}}_{\boldsymbol{imt}} = \arg\max \sum_{k \in \mathcal{K}_{imt}} (\tilde{p}_k - c_k)\, s_k,$$

where $\mathcal{K}_{imt}$ is the set of flights offered by airline i in market m and period t. The first-order optimality conditions for each product are described in equation (14). In a Nash-Bertrand equilibrium, each price choice satisfies (14), forming a system of nonlinear equations with dimension equal to the number of products. To simplify notation, drop subscripts mt and index products in a given market from 1 to K. Then, we can represent the stacked first-order conditions as

$$\underbrace{\begin{bmatrix} s_1 \\ s_2 \\ \vdots \\ s_K \end{bmatrix}}_{S} + \underbrace{\left\{ \begin{bmatrix} \frac{\partial s_1}{\partial \tilde{p}_1} & \frac{\partial s_2}{\partial \tilde{p}_1} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_1} \\ \frac{\partial s_1}{\partial \tilde{p}_2} & \frac{\partial s_2}{\partial \tilde{p}_2} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial s_1}{\partial \tilde{p}_K} & \frac{\partial s_2}{\partial \tilde{p}_K} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_K} \end{bmatrix} \circ \begin{bmatrix} O_{11} & O_{21} & \cdots & O_{K1} \\ O_{12} & O_{22} & \cdots & O_{K2} \\ \vdots & \vdots & \ddots & \vdots \\ O_{K1} & O_{K2} & \cdots & O_{KK} \end{bmatrix} \right\}}_{J} \underbrace{\begin{bmatrix} \tilde{p}_1 - c_1 \\ \tilde{p}_2 - c_2 \\ \vdots \\ \tilde{p}_K - c_K \end{bmatrix}}_{\tilde{P}-C} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \tag{A-18}$$

where $\circ$ is the element-wise (Hadamard) product. Cells O_{ij} indicate product ownership and receive 1 if products i and j are offered by the same airline and 0 otherwise. Partial derivatives of shares with respect to prices are obtained from the demand equation (12). Own and cross-price derivatives are given by

$$\frac{\partial s_k}{\partial \tilde{p}_k} = \frac{\alpha}{\lambda}(1+r)\, s_k \left[(1-\lambda)\, s_{k|g} + \lambda s_k - 1\right] \tag{A-19}$$

$$\frac{\partial s_k}{\partial \tilde{p}_j} = \frac{\alpha}{\lambda}(1+r)\, s_j \left[(1-\lambda)\, s_{k|g} + \lambda s_k\right], \tag{A-20}$$

where the tax term follows from $\frac{\partial p_k}{\partial \tilde{p}_k} = (1+r)$. Note that

$$\frac{\partial s_k}{\partial \tilde{p}_j} = \frac{\alpha}{\lambda}(1+r)\left[(1-\lambda)\frac{s_j s_k}{1-s_0} + \lambda s_j s_k\right] = \frac{\alpha}{\lambda}(1+r)\, s_k \left[(1-\lambda)\, s_{j|g} + \lambda s_j\right] = \frac{\partial s_j}{\partial \tilde{p}_k}.$$

Hence, since all products are in the same nest, the matrix of partial derivatives is symmetric. Because the ownership matrix is symmetric, the resulting matrix J is also symmetric. Define an auxiliary matrix Γ as

$$\Gamma \equiv \begin{bmatrix} \frac{1-\lambda}{1-s_0}s_1 + \lambda s_1 \\ \vdots \\ \frac{1-\lambda}{1-s_0}s_K + \lambda s_K \end{bmatrix} = \left(\frac{1-\lambda}{1-s_0} + \lambda\right) S = \left(\frac{1-\lambda s_0}{1-s_0}\right) S,$$

where s_0 is the share of the outside option (not flying). Then, we can write the matrix of partial derivatives as

$$\begin{bmatrix} \frac{\partial s_1}{\partial \tilde{p}_1} & \frac{\partial s_2}{\partial \tilde{p}_1} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_1} \\ \frac{\partial s_1}{\partial \tilde{p}_2} & \frac{\partial s_2}{\partial \tilde{p}_2} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial s_1}{\partial \tilde{p}_K} & \frac{\partial s_2}{\partial \tilde{p}_K} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_K} \end{bmatrix} = \frac{\alpha}{\lambda}(1+r) S\Gamma' - \frac{\alpha}{\lambda}(1+r) \begin{bmatrix} s_1 & 0 & \cdots & 0 \\ 0 & s_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & s_K \end{bmatrix}.$$

Thus, matrix J can be constructed using the following expression

$$J = \frac{\alpha}{\lambda}(1+r) \left\{ \left(\frac{1-\lambda s_0}{1-s_0}\right) SS' - \begin{bmatrix} s_1 & 0 & \cdots & 0 \\ 0 & s_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & s_K \end{bmatrix} \right\} \circ O, \tag{A-21}$$

where O indicates the ownership matrix.

Operating costs. Using (A-18), baseline operating costs are the solution to the following equation

$$\widehat{\boldsymbol{C}}_{\boldsymbol{mt}} = \tilde{\boldsymbol{P}}_{\boldsymbol{mt}} + \widehat{J}_{mt}^{-1} \boldsymbol{S}_{\boldsymbol{mt}},$$

where vectors $\tilde{\boldsymbol{P}}_{\boldsymbol{mt}}$ and $\boldsymbol{S}_{\boldsymbol{mt}}$ are observed in the data; matrix $\widehat{J}_{mt}$ is calculated using (A-21), so it depends on $\boldsymbol{S}_{\boldsymbol{mt}}$ and estimated parameters $\hat{\alpha}$ and $\hat{\lambda}$.

Counterfactual equilibria. As described in section 4.3, the baseline marginal operating cost per passenger is given by $c_k = \tilde{c}_k + (w_t + \tau_0) f_k$, where $\tilde{c}_k$ is the constant marginal cost per passenger excluding fuel costs, w_k is the jet fuel cost per gallon, and f_k is the volumetric fuel consumption per passenger. Hence, for a tax change from τ_0 to some tax level τ, operating cost increases by $\Delta c_k = f_k (\tau - \tau_0)$.

Let $\boldsymbol{C_{mt}}(\tau)$ denote the vector of operating costs after a tax change from the baseline. Then, the new Nash-Bertrand equilibrium satisfies

$$\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}} = -\widehat{J}_{mt}^{-1} \boldsymbol{S_{mt}} \left(\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}} \right) + \widehat{\boldsymbol{C}}_{\boldsymbol{mt}}(\tau), \tag{A-22}$$

where $\widehat{J}_{mt,\tau}^{-1}$ is itself a function of $\boldsymbol{S_{mt}}$, which in turn is a function of the pre-tax price vector $\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}}$. Since the right-hand side of (A-22) also depends on $\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}}$, the equilibrium is characterized by a fixed point problem.

I rewrite equation (A-22) as a root-finding problem, $\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}} + \widehat{J}_{mt}^{-1} \boldsymbol{S_{mt}} \left(\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}} \right) - \widehat{\boldsymbol{C}}_{\boldsymbol{mt}}(\tau) = 0$. To solve it, I use the trust region method with auto-scale and finite-difference approximated Jacobians (Nocedal & Wright, 2006). For each market and period, I solve for the equilibrium pre-tax price vectors over a grid of tax levels. This tax level grid ranges from $-\tau_0$ (removing the existing tax) to \$150/ton CO_2. The grid is adaptative and sets (i) intervals of \$1 around \$0 and \$108 to increase the precision of marginal changes from the baseline and the optimal tax level; (ii) intervals equivalent to a tenth of a cent change in the jet fuel tax (about 0.07 in the carbon tax) around \$61, to increase the precision of revenue-neutral tax effects; and (iii) intervals of \$10 elsewhere. To speed up the solver, I parallelize solutions across markets but solve them sequentially in the grid, using the previous closest available solution as the initial guess (or the observed price vector for nodes neighboring the baseline). With equilibrium prices, I calculate equilibrium shares and quantities, which are then used to calculate consumer surplus, profits, tax revenues, emissions, and damages.

A.4 GMM estimation with high-dimensional fixed effects

As discussed in section 6.2, estimating GMM with high-dimensional fixed effects and nonlinear parameters can create technical challenges. On the one hand, nonlinear parameters hinder the possibility of performing within transformation to eliminate fixed effects. On the other hand, leaving fixed effects as dummy variables increases memory and CPU requirements and may cause numerical instability when solving the problem with a large number of moment conditions. I address these issues by adapting the method proposed by Conlon and Gortmaker (2020), which I summarize below.

For a shorthand notation, let $\theta = [\alpha, \lambda]$ represent the vector of nonlinear parameters. Also, let $\theta_n = [\alpha_n, \lambda_n]$ represent the value of such parameters in the n-th iteration of the algorithm. We start with an initial guess θ_0. In this paper, I define the initial guess to be the parameters estimated using two-stage least squares. Nevertheless, varying the initial guess within reasonable values does not change the final estimates.

Step 1: Concentrate out nonlinear parameters. Define variable Y_k^D by rewriting the estimating equation for demand (20) as

$$Y_k^D \equiv (\ln s_k - \ln s_0) + \alpha_n p_k - (1-\lambda_n) \ln s_{k|g} = X_k^D \beta^D + \delta_i + \delta_{ot} + \delta_{dt} + \xi_k, \qquad \text{(A-23)}$$

where fixed effects δ_i, δ_{ot}, and δ_{dt} are separated from the vector of product characteristics. These fixed effects represent airline, origin location-by-quarter, and destination location-by quarter, respectively. Similarly, define Y_k^S by rewriting the estimating equation for supply (21) as

$$Y_k^S \equiv \underbrace{p_k - \mu_k(\alpha_n, \lambda_n)}_{c_k} = \rho F_k + \beta_i^S \text{Ramp-to-ramp}_k + \gamma_{i,o} + \gamma_{i,c_1} + \gamma_{i,d} + \gamma_{i,c_2} + \gamma_t + \omega_k$$

$$Y_k^S = X_k^S \beta^S + \gamma_{i,o} + \gamma_{i,c_1} + \gamma_{i,d} + \gamma_{i,c_2} + \gamma_t + \omega_k, \qquad \text{(A-24)}$$

where the term $\mu_k(\alpha, \lambda)$ makes it explicit that the markup used to calculate costs is a function of nonlinear parameters α and λ.

Let X^D be the $N \times M_D$ matrix of product characteristics relevant for demand (not including price and share within nest), where N is the number of products (observations) in the data, and M_D is the number of product characteristics excluding fixed effects; in this paper, $M_D = 8$. Let β^D be the $M_D \times 1$ vector of linear coefficients associated with demand characteristics. Similarly, let X^S be the $N \times M_S$ matrix of product characteristics relevant for supply, where $M_S = 8$. Also, β^S is the $M_S \times 1$ vector of linear parameters associated with X^S. Here, β^S includes parameter ρ and airline-specific parameters β_i^S.

Step 2: Absorb fixed effects. With the rearrangement of estimating equations outlined above, fixed effect terms in equations (A-23) and (A-24) can be eliminated. Since there are multiple fixed effects in each equation, simple demeaning would not work. Instead, we can absorb these terms using the method of alternating projections, which iteratively demeans variables until convergence (Bauschke, Deutsch, Hundal, & Park, 2003). Such a procedure is analogous to repeatedly applying the Frisch-Waugh-Lovell theorem in order to partial out dummy variables. The resulting equations can be written as

$$\begin{aligned}\bar{Y}_k^D &= \bar{X}_k^D \beta^D + \bar{\xi}_k \\ \bar{Y}_k^S &= \bar{X}_k^S \beta^S + \bar{\omega}_k,\end{aligned}$$

where bars on top of letters indicate demeaned variables.

Step 3: Estimate linear parameters. The moment conditions in equation (22) can be rewritten as

$$\begin{bmatrix} E\left(\bar{Z}_k^D \bar{Y}_k^D - \bar{Z}_k^D \bar{X}_k^D \beta^D\right) \\ E\left(\bar{Z}_k^S \bar{Y}_k^S - \bar{Z}_k^S \bar{X}_k^S \beta^S\right) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

with its sample analogue expressed as

$$\underbrace{\frac{1}{N}\begin{bmatrix} \left(\bar{Z}^D\right)' & 0 \\ 0 & \left(\bar{Z}^S\right)' \end{bmatrix}\begin{bmatrix} \bar{Y}^D \\ \bar{Y}^S \end{bmatrix}}_{\tilde{Y}} - \underbrace{\frac{1}{N}\begin{bmatrix} \left(\bar{Z}^D\right)' \bar{X}^D & 0 \\ 0 & \left(\bar{Z}^S\right)' \bar{X}^S \end{bmatrix}}_{\tilde{X}}\begin{bmatrix} \beta^D \\ \beta^S \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

where Z^D and Z^S denote the demand and supply instrument matrices, as outlined in section 6.2.

Note that $\bar{Y}^D$ and $\bar{Y}^S$ are functions of θ. Then, for a given value θ_n, we can estimate the linear parameters using GMM with weight matrix W:

$$\begin{bmatrix} \beta^D\left(\theta_n\right) \\ \beta^S\left(\theta_n\right) \end{bmatrix} = \left[\tilde{X}' W \tilde{X}\right]^{-1} \tilde{X}' W \tilde{Y}\left(\theta_n\right). \tag{A-25}$$

Since β parameters are linear, this step can be estimated relatively quickly by using routines optimized for matrix computation.

Step 4: Estimate nonlinear parameters. The nonlinear parameter vector θ is estimated using a two-step efficient GMM estimation (Hayashi, 2000). The estimated parameter satisfies

$$\theta^* = \arg\min_{\theta} NG(\theta)' \hat{W} G(\theta),$$

where $G(\theta)$ is the vector of moment conditions evaluated at θ and $\hat{W}$ is the GMM weighting matrix. In the first step, $\hat{W}_1$ is calculated based on the clustered residuals of a two-stage least squares regression. The solution to the first step is used to calculate matrix $\hat{W}_2$, and the second step solves for the nonlinear parameters that minimize the GMM objective function with this updated matrix.

In this paper, I solve the numerical GMM minimization problem over θ using the BFGS algorithm (Nocedal & Wright, 2006). For each iteration of θ (outer loop), the algorithm repeats steps 1 to 3 described above to calculate β^D, β^S, and residuals (inner loop).

B Additional regression results and alternative estimators

This appendix complements section 6 of the paper by showing estimation results using alternative estimators and specifications.

B.1 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 some-times 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-1 shows the estimates for the demand equation using 2SLS. This table shows that demand estimates using 2SLS and GMM (table 3) are generally similar but slightly smaller in absolute value. To provide further evidence of instrument validity, table A-2 shows the first-stage regressions for each instrumented variable, as well as the respective F-statistics. Though conditional F-statistics are substantially smaller than regular F-statistics, especially for price, these values are well above 13.95, the critical value at 5% with 3 endogenous variables and 7 excluded instruments (using values from Stock and Yogo (2002), as suggested in Sanderson and Windmeijer (2016)).

Finally, table A-3 shows estimates of supply-side parameters. Here, estimates are slightly different than 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.2 Alternative specifications

Section 6.2 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-4 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 by half, while other key parameters are less drastically affected. Interestingly, this smaller α is closer to the value of 0.45 found in Pagoni and Psaraki-Kalouptsidi (2016), which estimate a model without a rich set of location fixed effects.

In this main specification, prices are included at level because it is assumed that the representative consumer has a constant marginal utility of income equal to α—a standard assumption used in nested logit estimation and in other papers in the aviation literature. Relaxing this assumption involves reformulating the welfare framework and requires data on individual income, which, to my knowledge, does not exist. Nevertheless, we can evaluate whether the choice of using prices at level instead of log-prices is influential to other key parameters in the model. Specification (3) in Table A-4 shows that this is not the case: when including the natural log of prices, estimates of other parameters are minimally affected.

Specifications (4) and (5) in Table A-4 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) nest 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.

C Additional details on the data set

This appendix complements section 5 and presents additional information on definitions and variables that are calculated using multiple sources.

Markets geography. Figure A-2 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-5 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 I observe 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 0.217 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-3 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-1] The

[A-1]The number of available seats of an aircraft model varies by plane, with some airlines using more "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).

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.

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-4, 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

dense configurations. The calculation in this graph uses the average number of seats across all planes.

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.

D Additional results of 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-5 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-5 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-5 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.

E Market-specific marginal abatement costs

The estimated value of \$211/ton CO_2 corresponds to the marginal aggregate abatement cost based on a uniform carbon tax applied to all markets. Nevertheless, market power and ticket prices vary across markets, so tax wedges and abatement costs are heterogeneous. Figure A-6 illustrates the distribution of baseline MACs. Panel (a) in this Figure displays a histogram of the abatement costs per market; panel (b) shows cumulative distributions for the case where markets have equal weights (unweighted) and for the case where markets are weighted by their emissions. These graphs confirm the intuition that while existing distortions vary substantially, MACs are high even among the markets with the lowest abatement costs: the minimum MAC is \$107/ton CO_2. Thus, under the benchmark SCC of \$50/ton CO_2, there are no markets where a positive carbon tax would increase welfare.

The heterogeneity in costs also indicates inefficiencies arising from the use of a uniform carbon tax, as indicated in Section 3.4. Even though the second-best uniform tax of \$107/ton CO_2 (under a high SCC) improves aggregate welfare, it does so by taxing markets where the MAC is above the SCC of \$300—at the baseline, about a quarter of the markets representing 23% of emissions have MACs greater than \$300. Conversely, this uniform tax also undertaxes markets with low abatement costs. Nevertheless, market-specific carbon taxes would be difficult to implement in practice, especially considering that the optimal levels are based on heterogeneous market power.

Further examination of the distribution of MACs across markets indicates that emissions with lower abatement costs are mostly from large markets connecting dense urban areas, where competition leads to lower markups, fares, and, thus, tax wedges. Figure A-7 shows a histogram of MACs for each market size quartile, with the first quartile containing the smallest markets. This figure indicates that the dispersion of market MACs increases with market size: the proportion of markets with MAC below \$300/ton CO_2 is higher in the third and fourth quartiles.

Though the first quartile has a higher count of markets with low MAC, these markets generate a small fraction of the emissions in the sector. Since carbon damages are proportional to total emissions, it is also important to evaluate the distribution of MACs over emissions. To do so, I weigh each market by its emissions. Figure A-8 compares the unweighted and emission-weighted distributions of each market size quartile. Panel (a) in this figure reproduces the patterns observed in Figure A-7, with the lowest quartiles having a larger proportion of markets with MACs below $300. Panel (b) in Figure A-8 paints a similar picture with two exceptions. First, we observe that the distribution of the second quartile approaches those of higher quantiles, especially in the range of $200 and above. Second, the fourth quartile stands out in the low MAC range, indicating that a disproportional amount of the emissions is in markets with low distortions. This could reflect the fact that more competition in markets serving large urban areas tends to drive markups and prices down, thus decreasing the distortions that contribute to a higher MAC.

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.

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.

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

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

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

Hayashi, F. (2000). Econometrics. 2000. *Princeton University Press. Section*, *1*, 60–69.

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

Nocedal, J., & Wright, S. (2006). *Numerical optimization.* Springer Science & Business Media.

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

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.

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.

Stock, J. H., & Yogo, M. (2002). *Testing for weak instruments in linear IV regression.* NBER Working Paper 0284.

Figure A-1: Passenger traffic in cities and metro areas included in the data set.

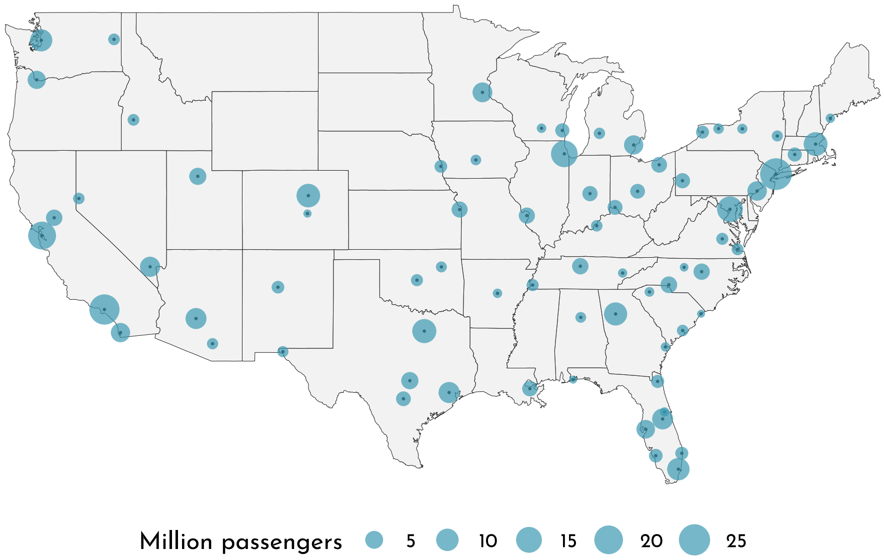

Source of data: US Bureau of Transportation Statistics. Note: traffic is measured in passengers enplaned in domestic flights from all airports within a city or metro area.

Figure A-2: Passenger traffic between cities and metro areas.

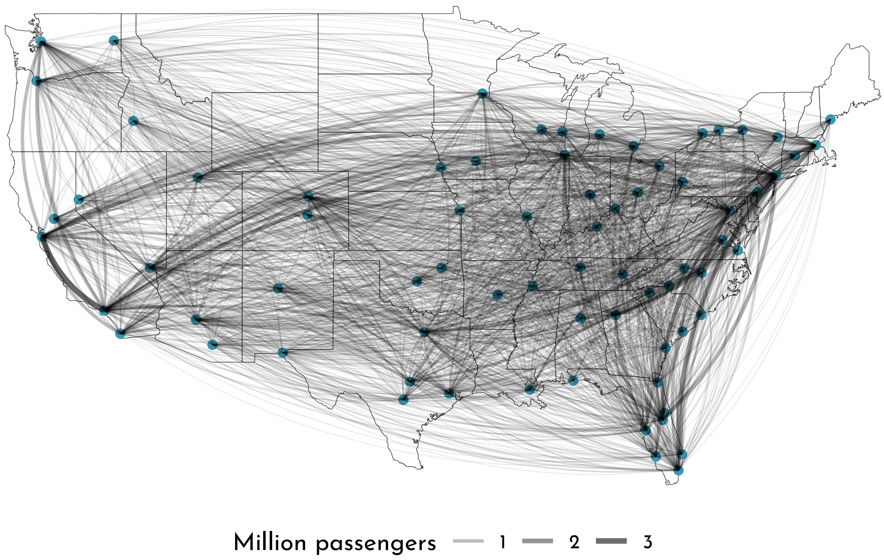

.

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-3: Fuel efficiency by stage length for the aircraft models most used.

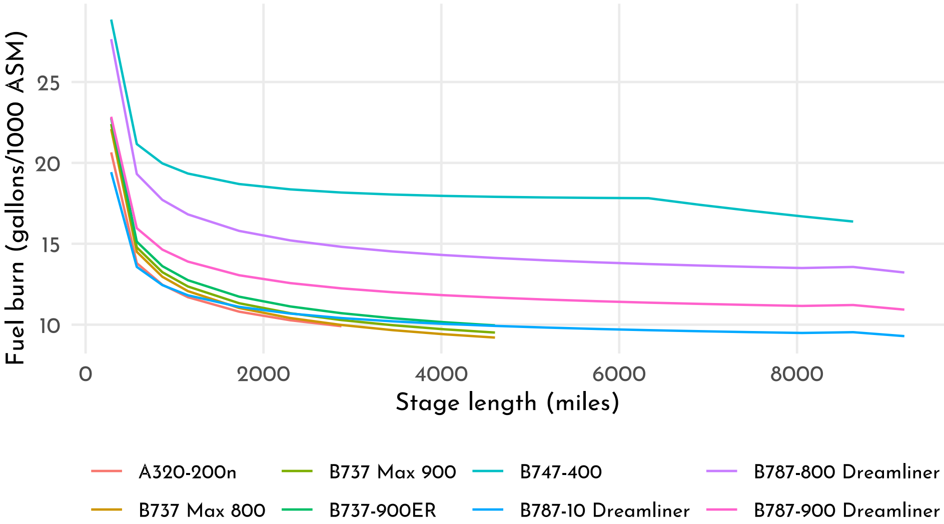

Notes: Fuel efficiency is measured in gallons of jet fuel per 1000 available seat-miles. The eight models shown in this graph correspond to those that offered the highest aggregate capacity (in available seat-miles) in 2018.

Figure A-4: Jet fuel and crude oil prices.

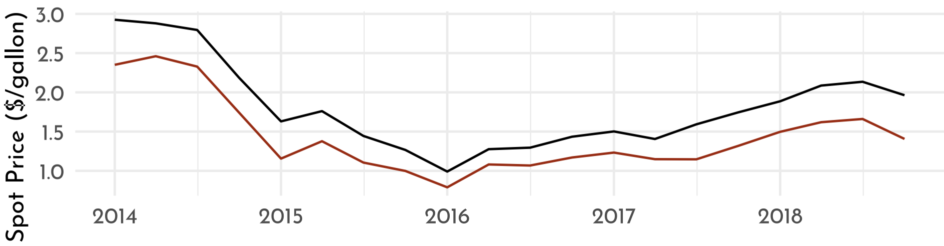

(a) Jet fuel and crude oil spot prices.

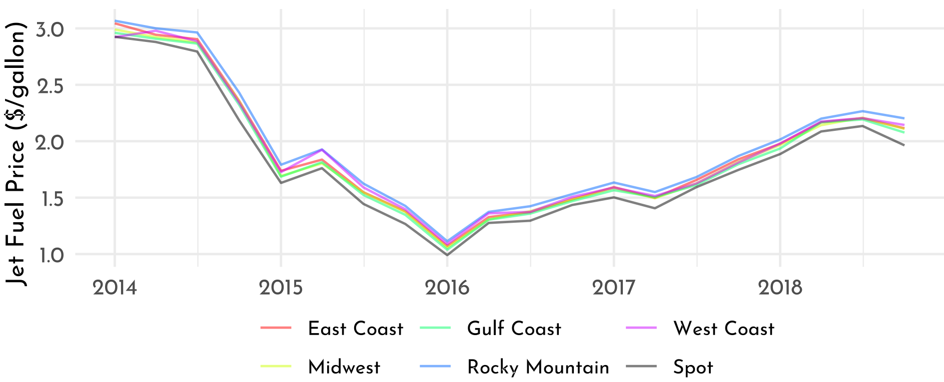

(b) Jet fuel spot and regional end-user sales prices.

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-5: Baseline characteristics of flights by carbon intensity quintiles.

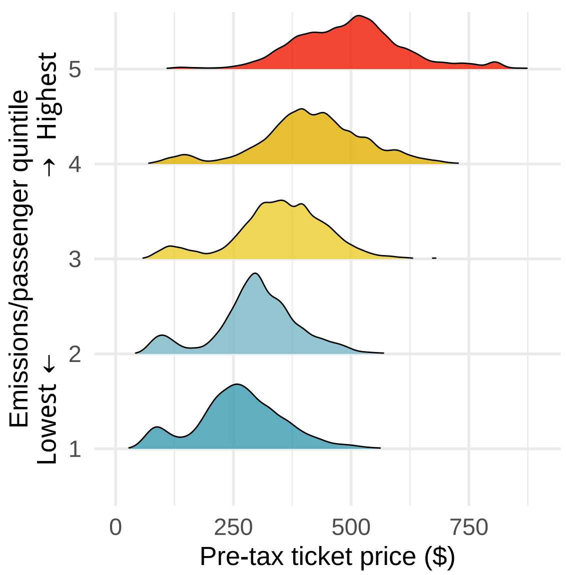

(a)

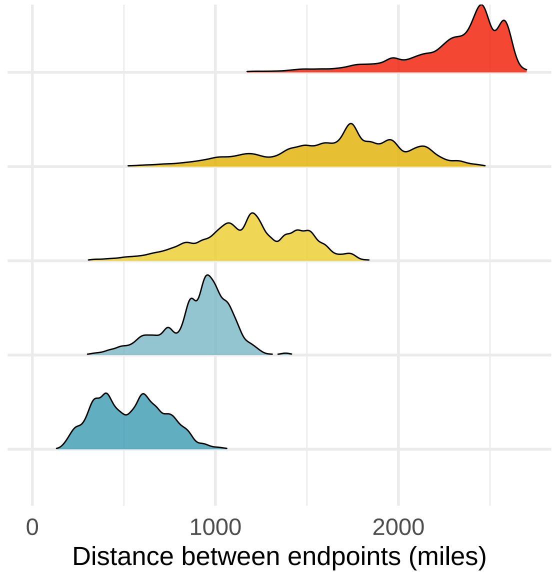

(b)

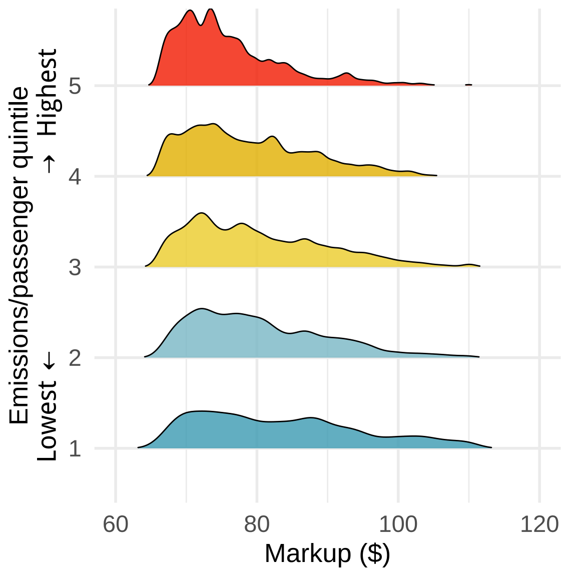

(c)

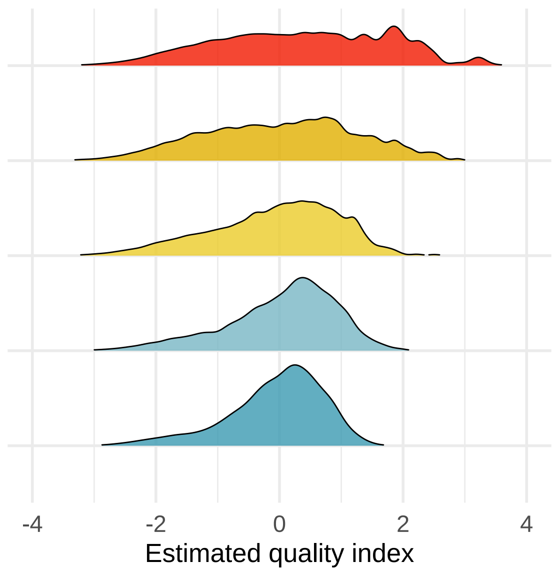

(d)

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-6: Distribution of baseline marginal abatement costs (MACs) across markets.

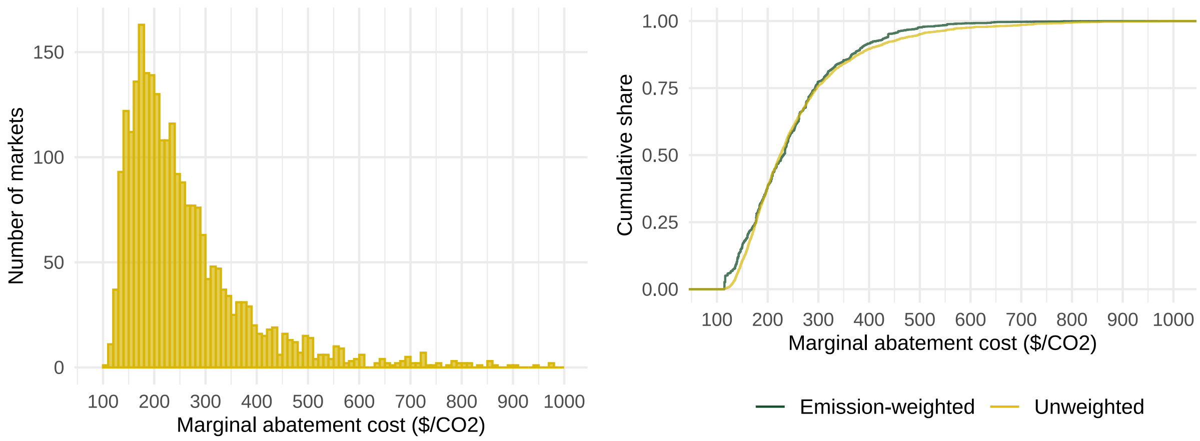

(a) Histogram.

(b) Cumulative distribution.

Notes: these calculations are based on the estimated model, as described in section 6.2. In panel (b), the *emission-weighted* distribution weights each market by its total emissions. This can be interpreted as the distribution of MACs with respect to total emissions, so that each point along the line indicates the share of emissions with a MAC at or below that level. In the *unweighted* distribution, shares are relative to the total number of markets.

Figure A-7: Histograms of baseline market-specific marginal abatement costs (MAC) for each market size quartile.

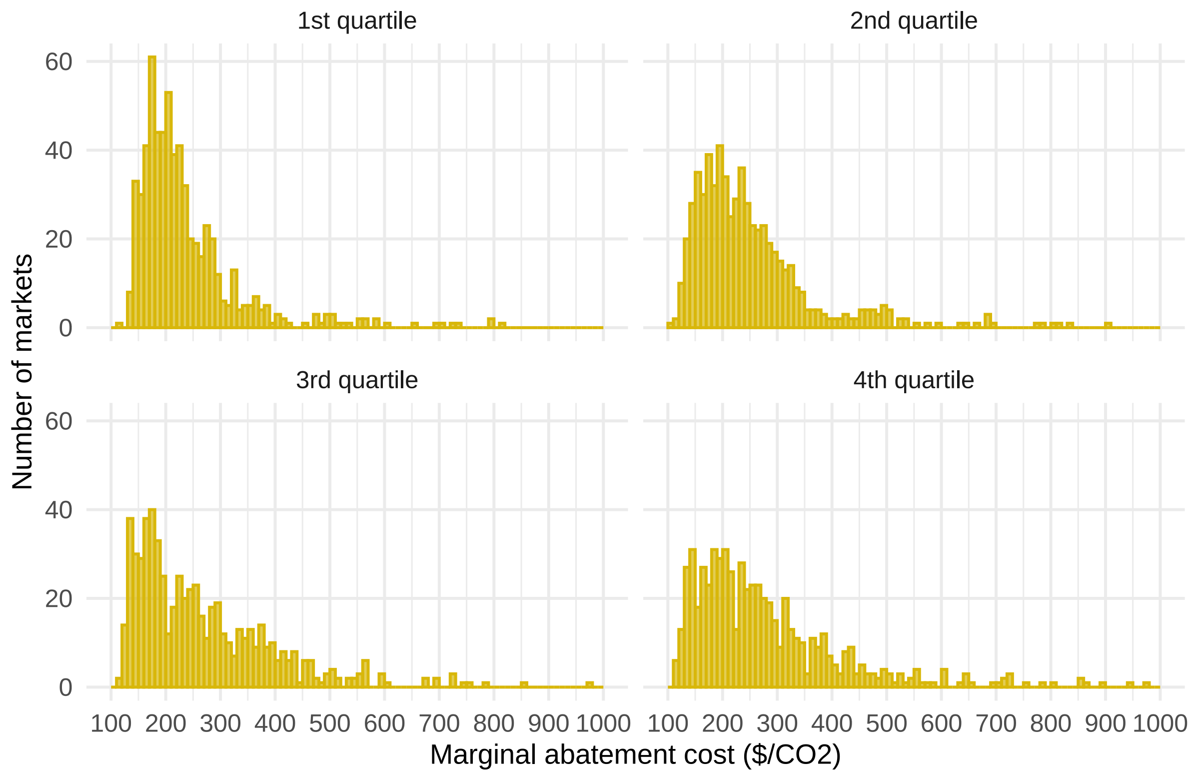

Notes: marginal abatement costs are defined in section 3.4. Smallest market sizes are in the first quartile.

Figure A-8: Cumulative distributions of baseline market-specific marginal abatement costs (MAC) for each market size quartile.

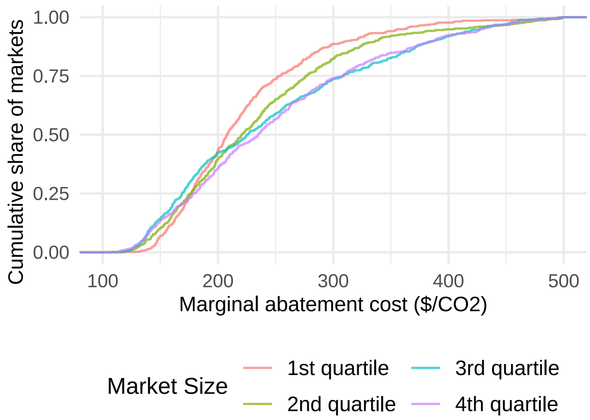

(a) Unweighted distribution

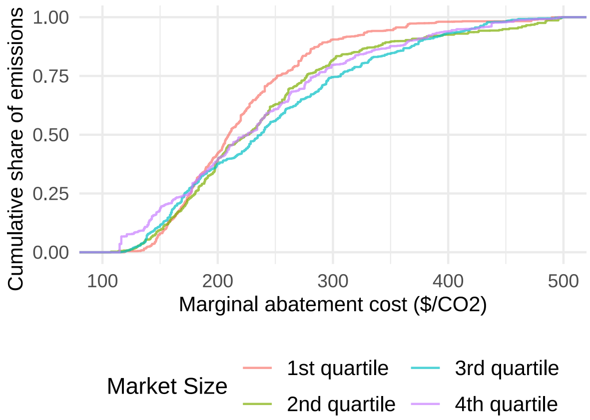

(b) Emission-weighted distribution

Notes: marginal abatement costs are defined in section 3.4. Smallest market sizes are in the first quartile. In panel (a), *unweighted* shares are relative to the total number of markets. In panel (b), *emission-weighted* distributions weight each market by its total emissions. These can be interpreted as the distributions of MACs with respect to total emissions, so that each point along the line indicates the share of emissions with a MAC at or below that level.

Table A-1: Estimates of demand parameters using 2SLS.

	$\ln(s_k/s_0)$
Price (\$100) [$-\alpha$]	-0.800
	(0.110)
ln(share within nest) [1-λ]	0.366
	(0.050)
Departures per week	0.036
	(0.002)
Number of stops	-0.848
	(0.073)
Market distance (100 mi.)	0.077
	(0.015)
Market distance squared	0.0001
	(0.0003)
Connection extra distance (100 mi.)	-0.089
	(0.009)
Connection extra distance squared	0.004
	(0.001)
Share of delayed departures	-0.603
	(0.094)
Destinations from origin	0.011
	(0.002)
Fixed effects	
Airline	Yes
Origin airport-by-quarter	Yes
Destination airport-by-quarter	Yes
Observations	267,967

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

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

	Dependent variable:		
	Price	ln(share within nest)	Departures per week
Number of stops	0.036	−1.179	−0.030
	(0.022)	(0.022)	(0.206)
Market distance (100 mi.)	0.068	0.016	−0.970
	(0.010)	(0.013)	(0.084)
Market distance squared	0.001	0.004	0.031
	(0.0002)	(0.0004)	(0.002)
Connection extra distance	−0.024	−0.152	−0.920
	(0.005)	(0.005)	(0.053)
Connection extra distance squared	0.001	0.010	0.076
	(0.0003)	(0.0004)	(0.004)
Share of delayed departures	−0.383	−0.416	−6.059
	(0.094)	(0.107)	(0.798)
Destinations from origin	0.018	−0.010	−0.086
	(0.001)	(0.001)	(0.006)
Airlines in market	−0.060	−0.129	−0.629
	(0.015)	(0.022)	(0.144)
Rivals' products in market	−0.0005	−0.002	0.008
	(0.0002)	(0.0004)	(0.002)
Rivals' % of nonstop flights	−0.602	−0.222	7.351
	(0.077)	(0.225)	(1.266)
Potential legacy entrants	0.074	−0.049	0.006
	(0.031)	(0.062)	(0.428)
Potential LCC entrants	−0.029	0.058	−0.857
	(0.012)	(0.017)	(0.113)
Fuel expenditure	0.003	−0.015	−0.066
	(0.001)	(0.001)	(0.010)
Compl. segment density	0.007	0.007	0.320
	(0.0004)	(0.0003)	(0.005)
Fixed effects			
Airline	Yes	Yes	Yes
Origin airport-by-quarter	Yes	Yes	Yes
Destination airport-by-quarter	Yes	Yes	Yes
Observations	267,967	267,967	267,967
F-statistic	317	957	353
Conditional F-statistic	23	62	62

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-3: Estimates of supply parameters using 2SLS.

	$\widehat{c}_k$
Fuel expenditure/avail. seat [ρ]	0.669
	(0.161)
Total ramp-to-ramp time (h)	0.156
	(0.017)
× *American*	−0.038
	(0.007)
× *Delta*	0.049
	(0.008)
× *United*	−0.014
	(0.008)
× *Alaska*	−0.020
	(0.033)
× *JetBlue*	−0.013
	(0.010)
× *Other low-cost*	−0.124
	(0.007)
Fixed effects	
Quarter	Yes
Each route airport-by-airline	Yes
Observations	267,967

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

Table A-4: Estimates of demand parameters under alternative specifications.

	Dependent variable:				
	$\ln(s_k/s_0)$				
	(1) OLS	(2) 2SLS No FE	(3) 2SLS Log prices	(4) 2SLS Airline nests	(5) 2SLS Flight type nests
Price (\$100) $[-\alpha]$	0.008	−0.411		0.101	−0.152
	(0.0004)	(0.085)		(0.082)	(0.025)
ln(Price)			−2.920		
			(0.453)		
ln(share within nest) [1-λ]	−0.040	0.443	0.381	0.411	1.075
	(0.002)	(0.021)	(0.041)	(0.034)	(0.017)
Departures per week	0.790	0.021	0.032	0.028	−0.001
	(0.008)	(0.001)	(0.002)	(0.002)	(0.001)
Number of stops	−0.346	−0.703	−0.803	−1.187	−0.310
	(0.014)	(0.031)	(0.061)	(0.018)	(0.016)
Market distance (100 mi.)	−0.007	0.037	0.067	0.008	0.003
	(0.009)	(0.013)	(0.014)	(0.009)	(0.003)
Market distance squared	−0.002	−0.001	−0.0003	−0.001	0.001
	(0.0004)	(0.0003)	(0.0003)	(0.0002)	(0.0001)
Conn. extra distance (100 mi.)	−0.027	−0.070	−0.093	−0.072	−0.009
	(0.002)	(0.005)	(0.007)	(0.007)	(0.002)
Conn. extra distance squared	0.001	0.003	0.004	0.001	0.001
	(0.0001)	(0.0003)	(0.0005)	(0.001)	(0.0001)
Share of delayed departures	−0.379	−0.799	−0.522	−0.682	0.176
	(0.054)	(0.079)	(0.088)	(0.075)	(0.027)
Destinations from origin	−0.001	−0.002	0.007	0.009	0.0001
	(0.0002)	(0.001)	(0.002)	(0.002)	(0.0005)
Fixed effects					
Airline	Yes	Yes	Yes	Yes	Yes
Origin airport-by-quarter	Yes	No	Yes	Yes	Yes
Destination airport-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

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

Table A-5: Airline groups and types in the dataset.

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