Pass-through, profits and the political economy of regulation

Felix Grey*
Faculty of Economics &
Energy Policy Research Group
Cambridge University
fg313@cam.ac.uk

Robert A. Ritz
Judge Business School &
Energy Policy Research Group
Cambridge University
rar36@cam.ac.uk

This version: January 2019

Abstract

Government regulation, such as the pricing of externalities, often raises firms' unit costs and its impact on their profits is important to its political economy. We introduce a new reduced-form model ("GLM") that nests existing models of imperfect competition. We show how firm-level cost pass-through is a sufficient statistic for the profit impact of regulation. We apply the GLM to carbon pricing for US airlines and find considerable intra-industry heterogeneity in pass-through. The GLM sidesteps estimation of a consumer demand system, firm mark-ups and conduct parameters. We derive a political-equilibrium emissions tax that incorporates firms lobbying a government "for sale".

Keywords: Cost pass-through, regulation, carbon pricing, airlines, political economy

JEL codes: D43 (imperfect competition), H23 (environmental externalities & taxes), L51 (regulation), L92 (air transport), Q54 (climate)

*Our thanks are due to Toke Aidt, Meredith Crowley, Natalia Fabra, Christos Genakos, Ken Gillingham, Pår Holmberg, Matt Kotchen, Hamish Low, Nathan Miller, Kamiar Mohaddes, Andrea Patacconi, Andrew Rhodes, and Donald Robertson for helpful discussions, to Severin Borenstein for providing airline data, and to Pradeep Venkatesh for excellent research assistance. We also thank seminar audiences at Cambridge, Dundee, Munich, Norwich Business School, Resources for the Future, Stockholm IFN, Yale and Wharton for their feedback. Felix Grey gratefully acknowledges financial support from the ESRC (Grant number ES/J500033/1); finally, we thank EPRG for supporting this research.

1 Introduction

We present a new approach to estimating the impact of regulation on firms' profits. It is based on a new reduced-form model of imperfect competition that unifies existing models using weaker assumptions. We apply the theory to understand the political economy of carbon pricing for the US airline industry, and quantify its winners and losers.

Government regulation often raises the production cost of regulated firms. In some cases, this is an explicit objective of regulation, for example, when it puts a price on an externality such as carbon emissions or uses an import tariff to protect domestic producers. In other cases, it is an inevitable consequence of a broader policy objective; examples are minimum wage legislation and capital requirements for banks. In addition to their effects on social welfare, such regulations have an important impact on the profits of the firms being regulated.[1]

This profit impact is critical to understanding the political economy of regulation. On the extensive margin, regulation that substantially lowers an industry's profitability is often unlikely to be introduced. On the intensive margin, firms may lobby the government to influence the equilibrium level of the regulation enacted. The profit impact is also important, for obvious reasons, to the shareholders of any regulated firm. For instance, major central banks are now among those warning institutional investors about the risks to asset values arising from climate-change policy (Carney 2015).

Estimating this firm-level profit impact is, however, not straightforward. Regulation raises the costs of a regulated firm and may also affect, to different degrees, the costs of its competitors. In general, its profit impact will depend on the firm's own production technology, the structure of demand, and its rivals' responses. The last factor is particularly problematic because modeling it may require information on the identities of all firms, each of their production technologies, the nature of product differentiation, what variables the firms compete on, how competitive or collusive the market is, and so forth. Our aim here is to present an approach that radically simplifies this problem.

The first half of this paper introduces our "generalized linear model" of competition (GLM), in the spirit of an "aim to build the theory in such a way as to focus attention on those predictions which are robust across a range of model specifications which are deemed reasonable" (Sutton 2007, pp. 2305-2306). The GLM makes weaker assumptions than typical models of imperfect competition. It assumes that (only) firm i is a cost-minimizer that takes input prices as given and operates a technology with linear production costs. The core assumption is that firm i follows a linear product-market strategy; in standard models, this corresponds to a linear supply schedule as implied by its first-order condition.

[1]See Draca, Machin & Van Reenen (2011) for a recent empirical study of the profit impacts of minimum wage legislation using UK data. In a distortion-free Modigliani-Miller world, the profit impact of bank capital regulation is zero; in practice, banking is characterized by imperfect competition alongside other market failures such as asymmetric information (see, e.g., Vives 2016).

There are no assumptions on the consumer demand system, no assumptions on the technologies and strategies of firm i's rivals, and no particular notion of "equilibrium".[2] The GLM also allows firm i to reduce its exposure to the regulated factor: under market-based environmental regulation, this is switching to cleaner inputs; faced with minimum-wage legislation, it is using less labour-intensive processes.

We use the GLM to characterize the impact of regulation that raises firm i's unit cost and affects those of its rivals in an arbitrary way. In doing so, we allow for "complete" regulation that covers an entire industry as well as "incomplete" regulation that exempts a subset of firms. We show how *firm-level* cost pass-through, i.e., the fraction of i's cost increase that is passed onto i's price, is a sufficient statistic for the profit impact of regulation.[3] That is, *all* relevant information on i's demand and supply conditions is contained in this single metric. We show that higher pass-through implies a more favourable profit impact; a firm's profit falls with tighter regulation if and only if its pass-through is below 100%. Up to this point, we treat pass-through as a parameter; it can be endogenized by assuming a particular mode of competition or estimated empirically.

To see the idea underlying the GLM, consider firm i which competes à la Cournot in a differentiated-products market, with marginal cost $MC_i = c_i + \tau$ and a demand curve $p_i = \alpha - \beta x_i - \delta \sum_{j \neq i} x_j$ (with $\delta \leq \beta$). Make no assumptions on its rivals' technologies or strategies. Firm i's first-order condition for profit-maximization implies a linear supply schedule $x_i = \frac{1}{\beta}(p_i - c_i - \tau)$. Now suppose that regulation τ raises i's own marginal cost by $d\tau$ and those of its rivals in an arbitrary way. How does this affect i's profits? By construction, i's pass-through rate $(dp_i/d\tau)/(dMC_i/d\tau)$ captures the impact on its profit margin $(p_i - MC_i)$. Moreover, due to the linear supply schedule, the change in its sales x_i is proportional to this pass-through rate. Rivals' cost shocks and competitive responses matter *only* insofar as they affect i's price—but this is precisely what is captured in i's pass-through rate. We show how to derive i's profit impact in a way that does not require knowledge of the demand parameters (α, β, δ) or of i's other costs c_i.

This basic logic extends to a rich class of oligopoly models. The GLM's structure nests, among others: Cournot-Nash, Stackelberg and conjectural-variation models (with linear demand); Bertrand and Cournot models with linear differentiated products; two-stage models such as Allaz & Vila (1993)'s model with forward contracting; a linear-symmetric version of the supply function equilibrium (Klemperer & Meyer 1989); behavioural theories of competition such as Al-Najjar, Baliga & Besanko (2008)'s model with sunk cost bias; and models with common ownership of firms (O'Brien & Salop 2000) which fea-

[2] An implication is that we can leave open (i) how many competing firms/products there are in the market; (ii) the extent to which firms' products are substitutes or complements in demand; (iii) the extent to which firms' products are strategic substitutes or strategic complements.

[3] We differ from the "firm-specific" rate of pass-through used in merger analysis (Ashenfelter, Ashmore, Baker & McKernan 1998); this asks how much of the cost saving achieved (only) by the merging firms is passed onto consumers. We also differ from a firm-specific pass-through rate based on a cost shock always incurred only by that single firm. Our setting allows for arbitrary cost changes across firms.

ture prominently in the current debate on the competitive impacts of institutional stock ownership (Azar, Schmalz & Tecu 2018).

While it is intuitive, the critical role played by pass-through is also far from obvious. In influential work, Weyl & Fabinger (2013) show, in a general class of symmetric oligopoly models, that a *market-wide* rate of cost pass-through is a useful tool to understand market performance. As Miller, Osborne & Sheu (2017) write: "the effect on producer surplus [of a market-wide cost shock] depends on pass-through and a conduct parameter that equals the multiplicative product of firm margins and the elasticity of market demand". By contrast, within the GLM, the firm-level profit impact depends *solely* on pass-through—no additional information about conduct parameter(s) is needed. This simplification of incidence analysis is the primary attraction of the GLM. On one hand, the GLM makes heavy use of the assumption that firm i employs a linear product-market strategy; on the other hand, it allows for near-arbitrary heterogeneity across firms.

Our approach also differs from the structural modelling often employed in empirical industrial organization (Bresnahan 1989; Berry, Levinsohn & Pakes 1995; Nevo 2001; Reiss & Wolak 2007; Einav & Levin 2010). By estimating a full set of primitives, structural models can be widely deployed to estimate merger impacts, the consumer value of new products, and so on. As Reiss and Wolak (2007, p. 39) put it: "The inferences that IO researchers draw about competition from price and quantity data rest on what researchers assume about demand, costs, and the nature of firms' unobservable strategic interactions".[4] In this paper, we sidestep estimation of a differentiated-products demand system, make no assumptions about the precise mode of strategic interaction, and show how firm-level pass-through is sufficient information to "close the model".[5] In our setting, pass-through therefore also captures the import of any departures from Nash and/or profit-maximizing behaviour. A drawback is that we are not able to perform counterfactual analysis.

The second half of the paper illustrates the utility of the theory. We estimate the profit impacts of (future) carbon pricing on the US aviation market. This setting is important in its own right: emissions from airline travel are projected grow well into the 21[st] century and economic regulation is likely as countries seek to implement internationally-agreed climate targets in a cost-effective manner. At our baseline carbon price of \$50 per ton of carbon dioxide, the annual "value" of US domestic airline emissions exceeds \$8 billion.[6]

[4]A structural IO model has three main ingredients: (i) the consumer demand system, often specified in a logit form; (ii) firms' production technologies, often relying, like us, on linearity of costs; and (iii) the mode of competition, often chosen as Bertrand-Nash. It then estimates own- and cross-price elasticities of demand as well as the competitiveness of the market via firms' first-order conditions.

[5]Put differently, consider a market with n firms. Structural IO modeling, in general, requires specification and estimation of n demand equations as well as n supply equations. With the GLM, we specify only i's supply curve and show how i's profit impact is fully captured by i's pass-through rate—which contains all relevant information about the remaining $2n-1$ model equations.

[6]Our baseline \$50/tCO_2 carbon price is illustrative but close to central estimates of the social cost of carbon (Nordhaus, 2017). Our estimated profit impacts generally scale linearly with the carbon price.

Like many other industries, aviation is characterized by important heterogeneities between firms—in terms of demand, costs and conduct. First, airlines' products are differentiated in terms of service quality, legroom, loyalty schemes, luggage allowances, and so on. Second, an airline's costs depend on its aircraft fleet (e.g., size, age, fuel efficiency) and its configuration, both of which vary widely across carriers. As a result, airlines often incur heterogeneous cost shocks—even when exposed to the same carbon price on the same route. Third, airlines operate different portfolios of routes (e.g., short-haul vs long-haul flights) and conduct across routes is heterogeneous at the airline-level (e.g., number of competing carriers, low-cost vs legacy carriers) and route-level (i.e., the same carrier competes differently on different routes). Hence it is difficult to know with confidence which specific model of competition is best suited to each individual route.

Leveraging the GLM, we estimate profit impacts while remaining agnostic about the precise mode of competition across routes. We have quarterly ticket price data for over 600 domestic US carrier-routes over the period 2002–2014 (yielding over 30,000 observations). We have detailed information on fuel costs at the carrier-route level; we use variation in fuel prices to estimate fuel cost pass-through, from which we predict carbon cost pass-through.[7] We estimate a separate regression for each carrier-route, thus allowing for heterogeneity both between carriers and within the product portfolio of each carrier. We then aggregate across routes to determine the overall profit impact of regulation for each airline.

Our results show considerable intra-industry heterogeneity in pass-through rates. The large legacy carriers (Alaskan, American, Delta, Hawaiian, United and US Airways) have pass-through of 55% on average across their routes. By contrast, the major low-cost carrier Southwest has a much higher pass-through rate of 148%.[8] Extrapolated to all US domestic routes, at a carbon price of \$50/tCO_2, we predict a profit *gain* for Southwest of \$1.5 billion and a combined profit loss of \$3 billion for legacy carriers.

What explains this pass-through heterogeneity? We decompose it into three parts: 60% of the difference arises due to different route portfolios: legacy carriers tend to fly longer routes which have lower pass-through; 20% is explained by Southwest using more fuel-efficient planes on the same routes; the final 20% can be attributed to differences in customer demand. A key implication is that, because of demand-side asymmetries, there is heterogeneous pass-through even for a cost shock that hits all firms equally.

[7] Given the current absence of carbon pricing at the federal level in the US, other recent work, like us, relies on temporal variation in fuel costs to proxy the impacts of future carbon pricing (e.g., Bushnell & Humber 2017; Miller, Osborne & Sheu 2017). For airlines, jet fuel costs are primarily driven by the global price of oil which an individual airline cannot influence.

[8] Pass-through above 100% is more readily interpretable in our setting than in other literature. First, we estimate pass-through *rates* (i.e., the \$-price response to a \$1 unit cost increase) rather than pass-through *elasticities* (i.e., the %-price response to a 1% unit cost increase) which are always lower. Second, our firm-specific cost pass-through reflects asymmetries in the cost shocks experienced by different firms; all else equal, a firm that experiences half the cost increase of a rival has pass-through twice as high.

We present three additional results on the determinants of pass-through, showing that: (i) legacy carrier pass-through is lower when facing *potential* competition from Southwest; (ii) legacy carrier pass-through is 15% higher, on average, during spells of bankruptcy; and (iii) higher volatility in jet fuel prices is associated with significantly lower pass-through, all else equal, both for Southwest and legacy carriers.

We close with a novel application to the political economy of regulation. The GLM allows us to unite two strands of literature: (1) an influential literature following Grossman and Helpman (1994) in which firms lobby the government over the strength of regulation and its policy is "for sale"; and (2) a classic literature following Buchanan (1969) on second-best emissions taxes in the presence of market power. The policymaker, cognizant of imperfect competition and under influence of lobbying, chooses her utility-maximizing level of regulation—which we call the "political equilibrium" tax.

We show in a unified model that equilibrium regulation is determined by cost pass-through.[9] The distortion away from the standard Pigouvian rule (emissions tax set at social marginal cost) is driven by industry profits: in (1), this determines the extent of lobbying; in (2), it effectively measures market power. The GLM tells us that firm-level pass-through pins down firm-level profit impacts—and so the industry-level analog is driven by a weighted average of pass-through rates across firms. This generalizes existing results on second-best emissions taxes (surveyed by Requate 2006) by allowing flexibility over the mode of competition and product differentiation, and clarifies the underlying economic intuition in terms of pass-through.

On the empirical side, our baseline estimate of the political-equilibrium tax at \$19/tCO_2 is less than half of a \$50/tCO_2 social cost of carbon (SCC). Average carbon cost pass-through is 78% so industry profits fall significantly with regulation. A decomposition shows that, of the \$31 shortfall, \$28 is due to the Buchanan market-power distortion while the remaining \$3 is due to Grossman-Helpman lobbying. The gains from weaker regulation accrue to the large legacy carriers. Our findings may help explain why aviation has, so far, been a climate laggard.

Plan for the paper. Section 2 explains how this paper contributes to the literature. Section 3 sets out the GLM, relates it to existing oligopoly models, and derives our main result on firm-level pass-through as a sufficient statistic. Section 4 begins with brief background on climate-change policy for aviation, and then presents our empirical analysis of carbon pricing for the US airline industry. Section 5 contains our application to the "political equilibrium" level of regulation. Section 6 concludes.

Appendix A presents several theoretical extensions to the GLM: (1) a simple model of second-degree price discrimination; (2) a linear multiproduct oligopoly based on the

[9] Conducting such welfare analysis requires a few more assumptions beyond those of the GLM, notably that consumers are utility-maximizers. The policymaker does not have access to other instruments, such as a price control, to directly address the market-power distortion.

"upgrades approach" (Johnson & Myatt 2003, 2006); (3) emissions abatement via end-of-pipe technology, B has additional information on our airline data, C gives further details on our empirical results, D contains extensive empirical robustness checks, and E presents further results on the determinants of airline cost pass-through.

2 Related literature

This paper contributes to three main strands of literature. First, a rich literature has estimated cost pass-through in response to a variety of cost shocks, including excise taxes, input prices, and exchange rates. Empirical work typically reports a single rate of cost pass-through at the market-level. Depending on the detailed context, it finds evidence of "incomplete" pass-through below 100% (e.g., De Loecker, Goldberg, Khandelwal & Pavcnik 2016), "complete" 100% pass-through (e.g., Fabra & Reguant 2014) as well as pass-through above 100% (e.g., Miller, Osborne & Sheu 2017).[10]

In this paper, we show the value of shifting attention to how pass-through behaves at the level of an individual firm. While prior work has emphasized *inter-industry* heterogeneity in pass-through due to differences in competition, demand and technology (Ganapati, Shapiro & Walker 2017), our empirical results highlight *intra-industry* pass-through heterogeneity. We show, in the context of airline competition, that pass-through heterogeneity is driven by asymmetries in the cost shocks experienced by individual firms as well as asymmetries in demand and conduct.[11] In particular, we find that firms can have different pass-through rates even for a uniform cost shock.

Market-wide pass-through has recently been actively used as a "bridge" between structural and reduced-form models (Weyl & Fabinger 2013; Atkin and Donaldson 2015; Bergquist 2017; Miller, Osborne & Sheu 2017). As explained above, the linearity of the GLM further simplifies this incidence analysis by showing how firm-level pass-through *alone* then becomes a sufficient statistic.[12]

Second, our paper adds to a growing environmental-economics literature that studies the impacts of emissions pricing on industry. This literature has so far mostly focused on markets with limited product differentiation: electricity and heavy industry such as cement and steel. A key theme is that the profit impacts of carbon pricing at the

[10]See Weyl & Fabinger (2013) for a useful discussion of the diverse set of empirical pass-through results. In earlier theoretical work, Bulow & Pfleiderer (1983) derive pass-through for a monopolist while Anderson, de Palma & Kreider (2001) generalize this analysis to different oligopolistic environments.

[11]Kim (2018) finds important asymmetries in pass-through rates in electricity generation where firms' production technologies are differentially exposed to changes in the (input) price of natural gas.

[12]Pass-through also plays a central role in the analysis of cartel damages. The "passing on" defense states that a plaintiff is not harmed by an upstream cartel overcharge to the extent that it is able to pass this onto its own downstream customers. Verboven & van Dijk (2009) highlight the additional output-contraction effect that arises due to pass-through and further reduces a plaintiff's profits. Our work shows that, for the GLM class of oligopoly models, the output-contraction effect is negative for any firm that does not fully pass on the cartel overcharge. Thanks to Natalia Fabra for this pointer.

industry-level are typically modest; this has been found using several different modeling approaches: general equilibrium (Bovenberg & Goulder 2005), Cournot-style oligopoly (Hepburn, Quah & Ritz 2013) and event study (Bushnell, Chyong & Mansur 2014). A key driver is that a substantial portion of carbon cost is passed through to consumers in the form of higher product prices. An important implication is that only a small fraction of allowances in a cap-and-trade scheme—typically covering no more than 20-30% of pre-regulation emissions—needs to be freely allocated in order to preserve industry profits (Hepburn, Quah & Ritz 2013). Any higher free allocation of emissions permits leads to "windfall profits" such that an industry actually *benefits* from carbon regulation.[13]

Our paper extends this literature both in terms of the modeling approach and empirical findings. Our GLM-based analysis extends the theory to richer modes of imperfect competition, incorporating differentiated products and firm heterogeneity. Our empirical findings for airlines also differ from previous literature: because of the large losses incurred by legacy carriers (owing to low pass-through), the industry-level profit impact of carbon pricing is more significant. We find that the allowance allocation needed to preserve industry-wide profits is 43%—considerably higher than in prior literature. (For legacy carriers only, the profit-neutral allocation is close to 100%.) Thus, our results suggest that the political economy of carbon pricing for differentiated-products industries may be substantially different from the markets such as electricity and cement on which the literature has focused to date.[14]

Third, we contribute to the industrial-organization literature on competition in the airline industry. This literature has been primarily concerned with estimating the competitiveness of the industry and issues of market structure (Brander & Zhang 1990; Kim & Sengal 1993; Berry & Jia 2010) and the role and impact of financial constraints (Busse 2002; Borenstein 2011); recent work has also highlighted differences between legacy and low-cost carriers and the special role played by Southwest (Goolsbee & Syverson 2008; Ciliberto & Tamer 2009). This paper finds new evidence showing that large heterogeneity across carriers exists in terms of pass-through as well. We also provide the first, to our knowledge, formal economic assessment of future US climate regulation on this industry that highlights how different carriers have different incentives to influence this regulation.[15]

[13] Another strand of work examines the impacts of incomplete regulation in which only a subset of firms is subject to carbon pricing. This tends to worsen the profit impacts experienced by regulated firms and has important consequences for the efficiency and design of environmental regulation (Fowlie, Reguant & Ryan 2016). The GLM theory developed in this paper allows for such incomplete regulation, while our empirical application to airlines involves a setting in which regulation is complete.

[14] Bushnell, Chyong & Mansur (2014) do not have estimates for airlines because their April 2006 event study pre-dates the inclusion of aviation into the EU Emissions Trading Scheme (EU ETS). The period they study was still characterized by near-full free allocation of emissions permits (rather than auctioning), and this is a key factor driving their finding that firms benefit from a higher carbon price.

[15] Our empirical analysis uses an average ticket price across passengers for each carrier-route; our data are not rich enough to estimate the impacts of regulation on price dispersion (Borenstein & Rose 1994;

3 The generalized linear model of competition

This section introduces a simple reduced-form model which we call the "generalized linear model" of competition (GLM). We developed the GLM to respond to Sutton's (2007) call for (industrial-organization) economists to derive "predictions which are robust across a range of model specifications which are deemed reasonable." We first set out and discuss the key features of the model and place it in the context of existing oligopoly models. We then derive our main result on the profit impact of regulation within the GLM.

As discussed in the introduction, we have in mind a regulation that leads to an increase in regulated firms' unit costs. Examples are the government putting a price on an externality or imposing minimum wage legislation in the labour market. The following exposition applies to a range of different types of regulation. We place particular emphasis on the case of environmental regulation; in the following section we then apply the GLM to estimate the impacts of emissions pricing on the US airlines industry.

3.1 Setup of the GLM

Firm i competes in an industry, selling an output quantity x_i of its product at a price p_i. Let e_i be one of the inputs that i uses in production. Regulation imposes a cost τ on each unit of this input e_i. In the case of environmental regulation, the regulated factor corresponds to firm i's emissions (e.g., of carbon dioxide); it is standard in the literature to view emissions as an input to production (Baumol & Oates 1988). Regulation then corresponds to putting a price τ on the environmental externality.

The regulation may also apply to all ("complete regulation"), some or none of the other firms in the industry ("incomplete regulation"). More specifically, let $\phi_j \in \{0, 1\}$ be an indicator variable which equals 1 if firm j is subject to the regulation and equals 0 otherwise. Our setup has $\phi_i = 1$ for firm i but does not rely one any specific assumptions about the ϕ_js of its rivals ($j \neq i$).

In general, firm i's profits can be written as $\Pi_i = p_i x_i - C_i(x_i, e_i) - \tau e_i$, where $p_i x_i$ is its sales revenue and its total costs are made up of its operating costs $C_i(x_i, e_i)$ plus its regulatory costs τe_i.

The GLM makes four assumptions about the production technology and supply behaviour of firm i. These are taken to hold over some interval $\tau \in [\underline{\tau}, \overline{\tau}]$ (with $\underline{\tau} \geq 0$) of interest over which the extent of regulation varies:

A1. (*Input price-taking*) Firm i takes input prices, including the regulation τ, as given.

Gerardi & Shapiro 2009). Our extension of the GLM to second-degree price discrimination, where early buyers pay less for their tickets (see Appendix A.1) shows that the profit impact of regulation remains exactly as in our main analysis, with i's pass-through defined in terms of the change in i's *average* price.

A1 is a standard assumption which is appropriate for many forms of regulation, including a tax on emissions.[16,17]

A2. (*Cost-minimizing inputs*) Firm i chooses its inputs, including the regulated factor e_i, optimally so as to minimize its total costs $C_i(x_i, e_i) + \tau e_i$ of producing output x_i.

A2 is a canonical assumption in microeconomic theory. For market-based environmental regulation, it implies the textbook result that, at the optimum, the emissions price equals i's marginal cost of reducing emissions, that is, $-\frac{\partial}{\partial e_i} C_i(x_i, e_i) = \tau$. If regulation applies to multiple firms ($\phi_j = 1$ for at least one firm $j \neq i$) their marginal costs of emissions reductions are equalized, yielding the well-known cost efficiency property.

A3. (*Constant returns to scale*) Firm i's optimized total costs faced with regulation τ are linear in output $C_i(x_i, e_i) + \tau e_i = k_i(\tau)x_i$, with unit cost $k_i(\tau) = c_i(\tau) + \tau z_i(\tau)$ where $c_i(\tau)$ is its per-unit operating cost and $z_i \equiv e_i/x_i$ is its "regulatory intensity" (use of the regulated factor per unit of output).

A3 is a more substantive assumption but is also common in the literature, including in pass-through analysis (Bulow & Pleiderer 1983; Anderson, de Palma & Kreider 2001; Weyl & Fabinger 2013), in the literature on environmental regulation under imperfect competition (Requate 2006; Fowlie, Reguant & Ryan 2016; Miller, Osborne & Sheu 2017) and in the analysis of the profit impacts of a minimum wage (Draca, Machin & Van Reenen 2011). It rules out, at least over the range $\tau \in [\underline{\tau}, \overline{\tau}]$, the presence of (binding) capacity constraints.

Combining A1–A3, standard production theory shows that, in response to tighter regulation $\tau'' > \tau'$, $z_i(\tau'') \leq z_i(\tau')$ and $c_i(\tau'') \geq c_i(\tau')$. In other words, the firm reduces its use of the regulated input and instead uses more of other inputs; this saves on direct regulation-related costs (lower z_i) but incurs higher unit costs on other inputs (higher c_i).[18] For environmental regulation, this represents emissions abatement: a lower emissions intensity $z_i(\tau'') \leq z_i(\tau')$ comes at a per-unit abatement cost $[c_i(\tau'') - c_i(\tau')]$.[19] If such factor substitution is infeasible or unprofitable then $z_i(\tau'') = z_i(\tau')$.

[16] It is also a common assumption for a cap-and-trade system in which the market price of emissions is determined by way of an emissions cap. The reason is that an individual firm is typically relative to the size of the overall system, often because the system covers a wide range of sectors.

[17] For now, we treat the regulation τ as exogenously given. In Section 5 we relax this assumption using a variant of the model in which firms, by making political contributions, lobby the government over the extent of regulation—which then also endogenizes τ.

[18] We do not require any specific functional-form assumptions on the relationship between z_i and c_i.

[19] The GLM's abatement technology is consistent with standard properties from the environmental-economics literature. Write i's operating costs as $C_i(x_i, e_i) = c_i(\tau)x_i$, where emissions e_i are optimally chosen given output x_i; equivalently, the emissions intensity $z_i \equiv e_i/x_i$ is optimally chosen given output x_i. The first property is that emissions and output are *complements*: $\frac{\partial}{\partial x_i} C_i(x_i, e_i) = c_i(\tau)$ and so, given x_i, higher e_i implies higher z_i and hence lower c_i, that is, $\frac{\partial^2}{\partial x_i \partial e_i} C_i(x_i, e_i) < 0$. The second property is

A key implication is that, by the envelope theorem, $dk_i(\tau)/d\tau = z_i(\tau)$, that is, firm i's unit cost increase arising from a *small* tightening in regulation is given by its optimized regulatory intensity at that level of regulation. At the optimum, the increased costs due to input substitution are of second order. Therefore, if the extent of regulation rises from an initial level $\underline{\tau}$ to a higher $\overline{\tau}$, the corresponding increase in i's optimal unit cost equals $\Delta k_i(\overline{\tau}, \underline{\tau}) = \int_{s=\underline{\tau}}^{\overline{\tau}} z_i(s)ds$.

Remark 1 *While our exposition focuses on regulation that is effectively an* input *tax, the GLM nests an* output *tax as a special case where firm i's regulatory intensity per unit of output satisfies* $z_i(\tau) \equiv 1$ *for all* $\tau \in [\underline{\tau}, \overline{\tau}]$.

Remark 2 *The GLM can also apply to command-and-control regulation for which the government mandates a particular usage of inputs. An example is the mandatory blending of biofuels into petrol. In such cases, i's unit cost increase* $dk_i(\tau)/d\tau = z_i(\tau)$ *arises from a regulation* τ *that is not an input price (and* $z_i(\tau) = z_i$ *if no factor substitution is feasible).*

Remark 3 *To see another application, consider minimum wage regulation (following Ashenfelter & Smith 1979; Draca, Machin & Van Reenen 2011). Firm i's regulated factor is labour, denoted by* e_i. *Firm i takes the minimum wage, denoted by* τ, *as given (A1) and chooses the quantity of labour employed optimally to minimize its costs (A2). It operates a production technology for which the unit cost of output* $k_i(\tau) = c_i(\tau) + \tau z_i(\tau)$ *is constant (A3), where* $z_i \equiv e_i/x_i$ *is the labour intensity of its output. In response to a higher minimum wage* τ, *firm i may respond by using more capital-intensive processes, reducing its labour intensity* $z_i(\tau)$ *but raising other costs* $c_i(\tau)$.

The final assumption is the key feature of the GLM:

A4. (*Linear product market behaviour*) Firm i's product market behaviour satisfies $x_i(\tau) = \psi_i [p_i(\tau) - k_i(\tau)]$, where $\psi_i > 0$ is a constant and $[p_i(\tau) - k_i(\tau)] > 0$ is its profit margin.

A4 says that firm i behaves such that its output is in (fixed) proportion to the profit margin it achieves. Intuitively, it sells more or prices higher into a more attractive market: its supply curve is (linearly) upward-sloping. As we discuss below, this feature is shared by many existing models of imperfect competition with linear demand structures. The substantive restriction here is that the proportionality factor ψ_i does not vary with the

that abatement costs are convex:

$$\frac{\partial C_i}{\partial e_i} = \left[\frac{\partial C_i}{\partial z_i}\right] \frac{\partial z_i}{\partial e_i} = \left[\left(\frac{\partial c_i}{\partial \tau}\frac{\partial \tau}{\partial z_i}\right) x_i(\tau)\right] \frac{1}{x_i(\tau)} = \frac{\partial c_i}{\partial \tau} \Big/ \frac{\partial z_i}{\partial \tau} < 0,$$

and so $\frac{\partial^2 C_i}{\partial e_i^2} > 0$, for given x_i, requires $\frac{\partial}{\partial z_i}\left(\frac{\partial c_i}{\partial \tau} / \frac{\partial z_i}{\partial \tau}\right) > 0$. The latter property is consistent with the GLM—but it is also not required for our main results.

regulation τ. While regulation can shift firm i's supply schedule it does not effect the *slope* of that schedule.

The GLM makes *no* assumptions on the technology, behaviour or rationality of firm i's rivals. These firms need not be input price-takers (A1), need not choose inputs optimally (A2), have constant-returns technologies (A3) or employ a linear product-market strategy (A4). Regulation may also affect the *set* of i's rivals, that is, induce new entry and/or exit.

The GLM also makes no assumptions on the demand system in the industry or on the nature of consumer behaviour. An implication is that we can leave open (i) how many competing firms/products there are in the market; (ii) the extent to which other firms' products are substitutes or complements to i's; (iii) the extent to which firms' products are strategic substitutes or strategic complements to i's.

Hence the GLM has no equilibrium concept; it does not necessarily restrict attention to a Nash equilibrium (or some variation thereof). In this sense, it is much more general than standard models in which all firms are assumed to be Nash profit-maximizers. Moreover, A4 is consistent with "rule-of-thumb" behaviour by firm i that may itself not be profit-maximizing (though we do require cost-minimization as per A2).

3.2 Special cases of the GLM

To illustrate the scope of the GLM, we next set out some examples of well-known oligopoly models for which A4 is satisfied by *all* $n \geq 2$ firms in the industry. In these models, given A1–A3 and a linear demand structure, each firm's first-order for profit-maximization yields a linear supply schedule that takes the form of A4. It will be clear that the following listing is not exhaustive; other models with similar linear structures are also members of the GLM.

- *Cournot competition with a linear market demand curve* $p = \alpha - \beta \sum_i x_i$. It is easy to check that Nash behaviour implies $\psi_i = \beta^{-1}$ $(\forall i)$. Including a *firm-specific* conjectural variation $v_i \equiv \left(\sum_{j \neq i} dx_j \right) / dx_i$ (Bresnahan 1989) leads to $\psi_i = [\beta(1 + v_i)]^{-1}$, still consistent with the GLM. This also nests as a special case the linear Stackelberg model with multiple leaders and multiple followers (Daughety 1990).

- *Bertrand or Cournot competition with horizontal and/or vertical product differentiation.* Let firm i's demand $p_i = \alpha_i - x_i - \delta \sum_{j \neq i} x_j$, where $\delta \in (0, 1)$ is an inverse measure of horizontal differentiation and $\alpha_i \neq \alpha_j$ reflects vertical differentiation; this leads to $\psi_i = 1$ $(\forall i)$ for Cournot and correspondingly to $\psi_i = [1 + \delta(n - 2)] / [(1 - \delta)[1 + \delta(n - 1)]]$ $(\forall i)$ for Bertrand (Häckner 2000). The latter is independent of τ as long as the set of active firms and degree of product heterogeneity are unaffected. A richer "semi-linear" demand system $p_i = \alpha_i - \beta_i x_i - f_i(x_1, ..., x_{j \neq i}, ..., x_n)$

remains part of the GLM, for *any* cross-price effects implied by the function $f_i(\cdot)$.[20]

- *Two-stage model with forward contracting (Allaz & Vila 1993).* With linear demand $p = \alpha - \beta \sum_i x_i$, the subgame-perfect equilibrium features $\psi_i = (\beta/n)^{-1}$ $(\forall i)$ (Ritz 2014); this is again independent of τ as long as n is fixed.[21] This modelling approach is widely used in the study of liberalized electricity markets (e.g., Bushnell, Mansur & Saravia 2008). The identical equilibrium also arises in the well-known two-stage model of managerial delegation (Vickers 1985; Fershtman & Judd 1987).

- *A linear version of supply function equilibrium (Klemperer & Meyer 1989).* Demand is linear $p = 1 - \sum_i x_i$ and firm i has a linear supply schedule of the form $x_i = \sigma_i + \mu(p - k)$, where it chooses σ_i and firms' marginal costs are identical $k_i = k$ $(\forall i)$ (Menezes & Quiggin 2012). Even though the strategy space has an affine supply function, the symmetric equilibrium features A4 with $\psi_i = [1 + (n-1)\mu]$ $(\forall i)$.[22]

- *Competition with behavioural biases: misallocation of sunk costs (Al-Najjar, Baliga & Besanko 2008).* Firm i maximizes accounting profits which erroneously include some of its fixed costs, at a rate of $s_i > 0$ per unit of output. Firms learn about the impacts of costing in a "distortion game". With differentiated Bertrand competition with linear demand $D_i = a - bp_i + g\sum_{j \neq i} p_j$, the first-order conditions feature a constant $\psi_i > 0$ (identical across firms), where $s_i = s^* > 0$ in symmetric equilibrium.[23]

- *Linear competition with common ownership between firms (O'Brien and Salop 2000).* If the shareholders of firm i also own a fraction ω_i of its rival firm j, then the incentives of i's managers will be to maximize $\Pi_i + \omega_i \Pi_j$. The implications of such shareholder diversification have recently received attention for US airlines (Azar, Schmalz & Tecu 2018) and several other industries. With Cournot competition, linear demand, and assuming symmetry, this yields $\psi_i = [\beta(1+\omega)]^{-1}$ $(\forall i)$.

The GLM is more flexible than standard models along important dimensions. First, firms within the GLM may think they are playing a different game. Second, they may be using different choice variables (e.g. one firm chooses price and another firm chooses

[20] Similarly, it is not difficult to check that spatial models of competition, such as Hotelling and Salop, can also be consistent with the GLM as long as the distribution of consumer valuations is uniform (which generates linear demand, and, given A1–A3, hence also A4).

[21] To be more precise, this involves a three-stage game: (1) the emissions price is determined; (2) firms engage in forward contracting; and (3) firms compete in the product market à la Cournot.

[22] We have modified the setup of Menezes & Quiggin (2012) slightly, by dropping their normalization of θ_i in terms of production costs, without affecting any of the conclusions. The equilibrium price in our specification is $p^* = (1 + nk\,[1 + (n-1)\mu])\,/\,(1 + n\,[1 + (n-1)\mu])$, which tends to the Cournot solution as $\mu \to 0$ and to the Bertrand Paradox as $\mu \to \infty$.

[23] Specifically, using the results in Section 5 of Al-Najjar, Baliga & Besanko (2008), $\psi_i = [b/(2b + g)]\left[b\,(1 + g/[2b - (n-1)g]) - (n-1)g^2/[2b - (n-1)g]\right]$.

quantity). Third, the GLM, by allowing ψ_i to vary across firms, does not impose that a firm with a higher market share necessarily has a higher profit margin. This feature is "baked in" to many standard oligopoly models via the restriction $\psi_i = \psi$ $(\forall i)$.

Remark 4 *The GLM is distinct from classes of oligopoly models which are aggregative games (Corchón 1994; Acemoglu & Jensen 2013) or potential games (Monderer & Shapley 1996). Cournot-Nash competition with linear market demand* $p = \alpha - \beta\sum_i x_i$ *is an aggregative game, a potential game and also a member of the GLM. However, with differentiated-products demand for firm* i *of* $p_i = \alpha - \beta_i x_i - \sum_{j \neq i} \delta_{ij} x_j$*, it is no longer aggregative nor a potential game but still yields A4 and thus remains within the GLM.*

3.3 Profit-neutrality technique

We wish to quantify the impact of regulation τ on firm i's profits Π_i, and now use a simple modelling device to pin this down. Suppose that the extent of regulation rises from initial level $\underline{\tau}$ to a higher $\overline{\tau}$. Of particular interest are the special cases in which (i) a new regulation is introduced, that is, $\underline{\tau} \equiv 0$, and (ii) regulation is tightened by a small amount, that is, $(\overline{\tau} - \underline{\tau}) \to 0$.

Let $\Pi_i(\tau)$ denote firm i's profits as a function of regulation, and similarly $e_i(\tau)$ is the quantity of the regulated factor. Observe that $(\overline{\tau} - \underline{\tau})e_i(\underline{\tau})$ is the "static" cost of regulation associated with firm i's *initial* quantity of the regulated factor.

Now define a profit-neutrality factor γ_i for firm i such that:

$$\Pi_i(\overline{\tau}) + \gamma_i \left[(\overline{\tau} - \underline{\tau})e_i(\underline{\tau})\right] \equiv \Pi_i(\underline{\tau}). \tag{1}$$

In other words, γ_i is *constructed* such that firm i's profits under tighter regulation plus "compensation" of $\gamma_i \left[(\overline{\tau} - \underline{\tau})e_i(\underline{\tau})\right]$ are equal to its initial profits under weaker regulation. Hence the change in firm i's profits equals $\Delta\Pi_i(\overline{\tau}, \underline{\tau}) \equiv \Pi_i(\overline{\tau}) - \Pi_i(\underline{\tau}) = -\gamma_i \left[(\overline{\tau} - \underline{\tau})e_i(\underline{\tau})\right]$, so $\text{sign}\{-\Delta\Pi_i(\overline{\tau}, \underline{\tau})\} = \text{sign}\{\gamma_i(\overline{\tau}, \underline{\tau})\}$. If firm i does not respond to the tightening of regulation in any way, and its rivals do not respond either, then its profits simply fall as implied by its initial exposure, $\Delta\Pi_i(\overline{\tau}, \underline{\tau}) = -(\overline{\tau} - \underline{\tau})e_i(\underline{\tau})$ or, equivalently, $\gamma_i(\overline{\tau}, \underline{\tau}) = 1$.

The profit impact is therefore determined by three components: the extent of regulatory tightening $(\overline{\tau} - \underline{\tau})$, firm i's initial exposure $e_i(\underline{\tau})$, and the profit-neutrality factor $\gamma_i(\overline{\tau}, \underline{\tau})$. As is standard, we take the first component as given for now. The second component is typically available to policymakers and analysts dealing with regulation; for example, information on firms' historical emissions is collected in the run-up to environmental initiatives. The remainder of this section detemines the final component: firm i's profit-neutrality factor $\gamma_i(\overline{\tau}, \underline{\tau})$.

3.4 Main result on the profit impact of regulation

We next present our main result on the magnitude of $\gamma_i(\overline{\tau}, \underline{\tau})$ within the GLM. We will see that cost pass-through plays a central role for the profit impact of regulation. Define firm i's *marginal* rate of cost pass-through as:

$$\rho_i(\tau) \equiv \frac{dp_i(\tau)/d\tau}{dk_i(\tau)/d\tau}. \tag{2}$$

The denominator captures by how much i's optimal unit cost $k_i(\tau)$ responds to a small tightening in regulation; as explained above, given A1–A3, the envelope theorem implies $dk_i(\tau)/d\tau = z_i(\tau)$. The numerator captures by how much i's product price changes. This will, in general be driven by a combination of i's own cost increase and the cost increases and product-market behaviour of i's rivals (also depending on the extent to which they are regulated, i.e., the ϕ_js ($j \neq i$)).

Indeed, in standard oligopoly models, if *all* $n \geq 2$ firms are regulated, i's equilibrium price is a function of the marginal costs of all of its rivals, $p_i(k_1(\tau), ..., k_i(\tau), ...k_n(\tau))$. Therefore the price-response $dp_i(\tau)/d\tau$ also captures any relevant changes in these costs. (Note that the GLM does *not* imply that this pricing function $p_i(k_1(\tau), ..., k_i(\tau), ...k_n(\tau))$ is necessarily linear in its arguments. For example, with standard Nash differentiated-products competition with linear demand $p_i = \alpha_i - x_i - \delta \sum_{j \neq i} x_j$ the pricing function is indeed linear but for semi-linear demand $p_i = \alpha_i - \beta_i x_i - f_i(x_1, ..., x_{j\neq i}, ..., x_n)$ it is not.)

Our firm-level pass-through rate thus reflects the impact of regulation on *all* players, including any cost heterogeneity across firms. It is distinct from the market-wide pass-through rate considered in much of the pass-through literature (see especially Weyl & Fabinger 2013), for which each firm's cost rises by a uniform amount.[24]

Let $\overline{\rho}_i(\overline{\tau}, \underline{\tau}) \equiv \frac{1}{(\overline{\tau} - \underline{\tau})} \int_{s=\underline{\tau}}^{\overline{\tau}} [\rho_i(s)]\, ds$ denote i's *average* rate of cost pass-through over the interval $\tau \in [\underline{\tau}, \overline{\tau}]$. We thus obtain our main result:

Proposition 1 *In the GLM defined by A1–A4, the profit impact on firm i of regulation τ tightening from $\underline{\tau}$ to $\overline{\tau}$ satisfies $\Delta\Pi_i(\overline{\tau}, \underline{\tau}) \equiv -\gamma_i(\overline{\tau}, \underline{\tau})(\overline{\tau} - \underline{\tau}) e_i(\underline{\tau})$ where:*

(a) if $(\overline{\tau} - \underline{\tau})$ is small, then $\gamma_i(\overline{\tau}, \underline{\tau}) \simeq 2[1 - \overline{\rho}_i(\overline{\tau}, \underline{\tau})]$, where $\overline{\rho}_i(\overline{\tau}, \underline{\tau}) \simeq \rho_i(\underline{\tau})$;

(b) in general, $\gamma_i(\overline{\tau}, \underline{\tau}) \leq \max\{2[1 - \overline{\rho}_i(\overline{\tau}, \underline{\tau})], 0\}$.

Proof. The proof begins by deriving a general expression for the profit impact, and then shows parts (a) and (b) of the result. In general, using A1, firm i's optimum profit as a function of regulation τ is $\Pi_i(\tau) = p_i(\tau)x_i(\tau) - C_i(x_i(\tau), e_i(\tau)) - \tau e_i(\tau)$. Using A2 and A3, $[C_i(x_i(\tau), e_i(\tau)) + \tau e_i(\tau)] = k_i(\tau)x_i(\tau)$ so this simplifies to $\Pi_i(\tau) = [p_i(\tau) - k_i(\tau)]x_i(\tau)$.

[24] If both demand and costs are symmetric across firms, and A1–A4 apply to all firms, then our measure of pass-through typically coincides with market-wide pass-through.

Using A4, this becomes $\Pi_i(\tau) = \psi_i[p_i(\tau) - k_i(\tau)]^2$.
For part (a), note that the profit-neutrality factor from (1) can be written as:

$$\gamma_i(\overline{\tau}, \underline{\tau}) = \frac{1}{e_i(\underline{\tau})} \left[-\frac{\Pi_i(\overline{\tau}) - \Pi_i(\overline{\tau})}{(\overline{\tau} - \underline{\tau})} \right] = \frac{1}{e_i(\underline{\tau})} \left[-\frac{\Delta\Pi_i(\overline{\tau}, \underline{\tau})}{(\overline{\tau} - \underline{\tau})} \right].$$

Therefore, for small values of $(\overline{\tau} - \underline{\tau})$:

$$\gamma_i(\overline{\tau}, \underline{\tau})|_{\overline{\tau} \to \underline{\tau}} = \frac{1}{e_i(\underline{\tau})} \left(-\left.\frac{d\Pi_i(\tau)}{d\tau}\right|_{\tau = \underline{\tau}} \right).$$

Differentiation of the profit function $\Pi_i(\tau) = \psi_i[p_i(\tau) - k_i(\tau)]^2$ yields:

$$\left.\frac{d\Pi_i(\tau)}{d\tau}\right|_{\tau = \underline{\tau}} = 2\psi_i[p_i(\underline{\tau}) - k_i(\underline{\tau})] \left(\left.\frac{dp_i(\tau)}{d\tau}\right|_{\tau = \underline{\tau}} - \left.\frac{dk_i(\tau)}{d\tau}\right|_{\tau = \underline{\tau}} \right).$$

By the definition of (2), marginal pass-through $\rho_i(\underline{\tau}) = [dp_i(\tau)/d\tau]_{\tau = \underline{\tau}} \,/\, [dk_i(\tau)/d\tau]_{\tau = \underline{\tau}}$, and also by definition i's quantity of the regulated factor equals its output times its regulatory intensity, $e_i(\underline{\tau}) = z_i(\underline{\tau}) x_i(\underline{\tau})$. By A1–A3 and the envelope theorem, $[dk_i(\tau)/d\tau]_{\tau = \underline{\tau}} = z_i(\underline{\tau})$ and, by A4, $x_i(\underline{\tau}) = \psi_i[p_i(\underline{\tau}) - k_i(\underline{\tau})]$. So $[d\Pi_i/d\tau]_{\tau = \underline{\tau}} = 2e_i(\underline{\tau})\,[\rho_i(\underline{\tau}) - 1]$. Combining results, it follows that, for small values of $(\overline{\tau} - \underline{\tau})$:

$$\gamma_i(\overline{\tau}, \underline{\tau}) \simeq \gamma_i(\overline{\tau}, \underline{\tau})|_{\overline{\tau} \to \underline{\tau}} = 2\,[1 - \rho_i(\underline{\tau})]\,, \tag{3}$$

where, locally, marginal pass-through approximately equals average pass-through, $\rho_i(\underline{\tau}) \simeq \overline{\rho}_i(\overline{\tau}, \underline{\tau})$.
For part (b), the change in profits $\Delta\Pi_i(\overline{\tau}, \underline{\tau}) = \psi_i \left\{ [p_i(\overline{\tau}) - k_i(\overline{\tau})]^2 - [p_i(\underline{\tau}) - k_i(\underline{\tau})]^2 \right\}$. Defining $\Delta p_i(\overline{\tau}, \underline{\tau}) \equiv [p_i(\overline{\tau}) - p_i(\underline{\tau})]$ and $\Delta k_i(\overline{\tau}, \underline{\tau}) \equiv [k_i(\overline{\tau}) - k_i(\underline{\tau})]$ and expanding and simplifying yields:

$$\Delta\Pi_i(\overline{\tau}, \underline{\tau}) = \psi_i \left\{ 2\,[p_i(\underline{\tau}) - k_i(\underline{\tau})]\,[\Delta p_i(\overline{\tau}, \underline{\tau}) - \Delta k_i(\overline{\tau}, \underline{\tau})] + [\Delta p_i(\overline{\tau}, \underline{\tau}) - \Delta k_i(\overline{\tau}, \underline{\tau})]^2 \right\}. \tag{4}$$

Recalling that $dk_i(\tau)/d\tau = z_i(\tau)$, it follows that $\Delta k_i(\overline{\tau}, \underline{\tau}) = \int_{s=\underline{\tau}}^{\overline{\tau}} z_i(s)ds$. This in conjunction with (2) gives $\Delta p_i(\overline{\tau}, \underline{\tau}) = \int_{s=\underline{\tau}}^{\overline{\tau}} \rho_i(s) z_i(s) ds$ and so $[\Delta p_i(\overline{\tau}, \underline{\tau}) - \Delta k_i(\overline{\tau}, \underline{\tau})] = \int_{s=\underline{\tau}}^{\overline{\tau}} z_i(s)[\rho_i(s) - 1]ds$. Using these expressions and $e_i(\underline{\tau}) = z_i(\underline{\tau}) x_i(\underline{\tau})$ in the profit-neutrality factor from (1) gives:

$$\gamma_i(\overline{\tau}, \underline{\tau}) = \frac{\psi_i \left\{ 2\,[p_i(\underline{\tau}) - k_i(\underline{\tau})] \left[\int_{s=\underline{\tau}}^{\overline{\tau}} [1 - \rho_i(s)] z_i(s) ds \right] - \left[\int_{s=\underline{\tau}}^{\overline{\tau}} [1 - \rho_i(s)] z_i(s) ds \right]^2 \right\}}{(\overline{\tau} - \underline{\tau}) z_i(\underline{\tau}) x_i(\underline{\tau})}.$$

From A4, $x_i(\underline{\tau}) = \psi_i\,[p_i(\underline{\tau}) - k_i(\underline{\tau})]$ and some rearranging gives a general expression for the profit impact:

$$\gamma_i(\overline{\tau},\underline{\tau}) = 2\left[\frac{1}{(\overline{\tau}-\underline{\tau})}\int_{s=\underline{\tau}}^{\overline{\tau}}\frac{z_i(s)}{z_i(\underline{\tau})}[1-\rho_i(s)]ds\right] - \frac{(\overline{\tau}-\underline{\tau})z_i(\underline{\tau})}{[p_i(\underline{\tau})-k_i(\underline{\tau})]}\left[\frac{1}{(\overline{\tau}-\underline{\tau})}\int_{s=\underline{\tau}}^{\overline{\tau}}\frac{z_i(s)}{z_i(\underline{\tau})}[1-\rho_i(s)]ds\right]^2 \tag{5}$$

Observe that the profit-neutrality factor is bounded above according to:

$$\gamma_i(\overline{\tau},\underline{\tau}) \leq 2\left[\frac{1}{(\overline{\tau}-\underline{\tau})}\int_{s=\underline{\tau}}^{\overline{\tau}}\frac{z_i(s)}{z_i(\underline{\tau})}[1-\rho_i(s)]ds\right], \tag{6}$$

where $z_i(s) \leq z_i(\underline{\tau})$ for all $s \in [\underline{\tau},\overline{\tau}]$ given A1–A3. There are two cases. If $\rho_i(s) > 1$, then $[z_i(s)/z_i(\underline{\tau})]\,[1-\rho_i(s)] < 0$. If $\rho_i(s) \leq 1$, then $0 < [z_i(s)/z_i(\underline{\tau})]\,[1-\rho_i(s)] \leq [1-\rho_i(s)]$. Therefore, whatever the value of $\rho_i(s)$, $[z_i(\underline{\tau})/z_i(0)]\,[1-\rho_i(s)] \leq \max\{0,[1-\rho_i(s)]\}$. It follows that $\gamma_i(\overline{\tau},\underline{\tau}) \leq \max\{0, 2[1-\overline{\rho}_i(\overline{\tau},\underline{\tau})]\}$, where average pass-through $\overline{\rho}_i(\overline{\tau},\underline{\tau}) \equiv \frac{1}{(\overline{\tau}-\underline{\tau})}\int_{s=\underline{\tau}}^{\overline{\tau}}[\rho_i(s)]\,ds$, as claimed. ■

Proposition 1 gives a very simple expression that makes precise how firm i's rate of cost pass-through *alone* is a *sufficient statistic* for the profit impact of regulation. It holds across all models that are part of the GLM. Conditional on the extent of regulation and firm i's "historical" use of the regulated factor, firm-level pass-through is the only thing that matters: $\Delta\Pi_i(\overline{\tau},\underline{\tau}) \simeq 2[1-\overline{\rho}_i(\overline{\tau},\underline{\tau})](\overline{\tau}-\underline{\tau})e_i(\underline{\tau})$.

Firm i's pass-through rate captures *all* relevant information on the production technologies of i's rivals, the degree of product differentiation, what variables the firms compete on, how competitive or collusive the market is, any entry or exit by rivals, and so on. Whatever their other differences, if any two theories within the GLM imply identical pass-through for firm i, then they also imply an identical profit impact.

For a "small" regulatory tightening in part (a), this argument holds exactly; for a "large" regulation in part (b), the profit impact is bounded above by the same expression.

What is the intuition for the result? The profit impact is made up of two effects: that on i's profit margin and that on its sales. The first role of i's pass-through rate is that, by construction, it captures the impact of regulation on its own profit margin. Its second role is that, due to the linear supply schedule given by A4, the change in its sales is proportional to its pass-through rate. Rivals' cost shocks and competitive responses matter *only* insofar as they affect i's price—but this is precisely what is captured in i's pass-through rate. These two roles are what drives $\gamma_i(\overline{\tau},\underline{\tau}) \simeq 2[1-\overline{\rho}_i(\overline{\tau},\underline{\tau})]$.

Two corollaries are immediate. First, pass-through *signs* the profit impact: i's profits fall whenever pass-through is incomplete, $\text{sign}(\Delta\Pi_i) = \text{sign}(\overline{\rho}_i - 1)$. In such cases, the firm's profit margin shrinks and, by A4, it also experiences weaker sales. Conversely, with pass-through above 100%, the firm benefits from tighter regulation: both its profit margin and sales volume rise. Second, all else equal, a lower rate of pass-through implies that regulation has a worse profit impact.[25]

From an empirical perspective, the crucial feature of the result is that the profit impact $\Delta\Pi_i(\tau)$ is *independent* of the proportionality term ψ_i from A4. As we have seen, different theories of imperfect competition within the GLM differ in terms of the their implied ψ_i. Yet Proposition 1 tells us that this does not matter for the profit impact. The reason is scaling: by A4, the level of $x_i(\tau) = \psi_i [p_i(\tau) - k_i(\tau)]$ and the change $\Delta x_i(\tau) = \psi_i [\Delta p_i(\tau) - \Delta k_i(\tau)]$ are both proportional to ψ_i. But the corresponding use of the regulatory factor $e_i(\tau) = \psi_i z_i(\tau) [p_i(\tau) - k_i(\tau)]$ is also proportional to ψ_i. This means that the profit impact *per unit* of the regulatory factor $e_i(\underline{\tau})$, as incorporated into $\gamma_i(\overline{\tau}, \underline{\tau})$, does not depend on ψ_i. This is also one reason for why Proposition 1 applies without requiring information about own-price and cross-price elasticities of demand.

Of course, industry characteristics such as the degree of product differentiation are likely to affect the pass-through rate ρ_i—so they can certainly matter *indirectly* for $\Delta\Pi_i$. The point of Proposition 1, is that, even if they also affect i's supply behaviour, via A4's proportionality term ψ_i, this aspect is irrelevant for the profit impact (conditional on ρ_i).

While it is intuitive, this critical role played by pass-through is also far from obvious. Weyl & Fabinger (2013) present, for a general class of symmetric oligopoly models, a simple formula for the impact of a *market-wide* cost change on aggregate producer surplus (see also Atkin & Donaldson 2015; Miller, Osborne & Sheu 2017). The profit impact depends on a market-wide rate of pass-through as well as on a "conduct parameter" that incorporates the level of firms' profit margins and the price elasticity of market demand. By contrast, within the GLM, the firm-level profit impact depends *solely* on pass-through—no additional information about conduct parameter(s) is needed. This further simplification of incidence analysis is the primary attraction of the GLM. Compared to the existing literature, the GLM allows for near-arbitrary heterogeneity across firms but makes heavy use of the linear structure implied by A4.

One interpretation is that the GLM is "semi-parametric" or "partially specified" as a model but that pass-through information is, in effect, sufficient to "close the model". Consider a market with n firms selling differentiated products. Standard industrial-organization models specify n demand equations and, based on profit-maximization, then derive n supply equations. With the GLM, we specify only i's supply curve and show how i's profit impact is fully captured by i's pass-through rate—which contains *all* relevant information about the remaining $2n - 1$ model equations. In our setting, pass-through

[25] We here comment on the question of aggregation across firms in cases where the GLM is assumed to hold for *all* firms in an industry. If firm-level pass-through exceeds 100% for all firms, A4 implies that sales rise for all firms—which, in equilibrium, contradicts higher prices given that demands are downward-sloping. So, in an equilibrium context, given A1–A4, pass-through can exceed 100%, so that profits rise, for at most $n-1$ firms in the industry. An important point is that, in an empirical application, if demands are indeed downward-sloping, then this will be reflected in the observed data. Therefore it is not necessary to impose this as separate condition in an empirical pass-through analysis. Indeed, in our application to US airlines, we find that firm-level pass-through exceeds 100% only for a single player.

therefore also captures the import of departures from Nash and/or profit-maximizing behaviour—including any player "irrationality".

It is clear that the formula from Proposition 1 can still *approximately* hold even if A3 or A4 are violated, that is, if firm i faces a mildly non-linear demand curve or had a production technology with slightly increasing or decreasing returns to scale.[26] An important observation is that such modest departures from the GLM do not introduce a *systematic* bias to the results. For example, if A1–A3 hold and n symmetric firms play Cournot-Nash (Hepburn, Quah & Ritz 2013), then it is easy to check that the formula from Proposition 1(a) will overstate the true profit decline if the demand curve is concave but understate it with convex demand.

Up to this point, we have treated i's pass-through rate as a parameter. In general, it can depend on a wide range of factors, including the number of firms in the market, their strategies and the intensity of competition, the degree of product differentiation, and on how strongly i's cost rises relative to other firms. To progress further, there are three basic approaches. The first is to select a specific model of competition and derive the theoretical rate of pass-through. The second is to try to utilize existing estimates of cost pass-through from the empirical literature (given that these are typically market-level pass-through estimates, they would have to be converted into firm-level pass-through, perhaps again relying on guidance from a specific theory model.) The third is to combine the structural result from Proposition 1 with new firm-level estimates of pass-through. We pursue this last approach in the next section, for the US airline industry.

In Appendix A, we present three extensions of the GLM. The first is a simple model of oligopolistic second-degree price discrimination in which customers who buy earlier pay less—an important feature, for example, of the airline market. Pass-through is then defined in terms of the *average* price paid by consumers. The second is a linear version of the "upgrades approach" to multiproduct quality competition (Johnson & Myatt 2003, 2006). The third has firm i instead invest in an "end of pipe" abatement technology that cleans up its production *ex post*. We show that, in all three settings, Proposition 1(a) on firm-level cost pass-through as a sufficient statistic is preserved.

4 Empirical analysis of carbon cost pass-through for US airlines

This section illustrates the utility of the theory: Using the GLM, we estimate the profit impacts of (future) carbon pricing on the US aviation market. We begin with brief background on aviation and climate change policy. We then discuss our strategy to move

[26] The literature on merger analysis also finds that, at least in some cases, logit demand systems—a popular form of non-linear demand—give quite similar results to those under linear demand (see, e.g., Crooke, Froeb, Tschantz & Werden 1999; Miller, Remer, Ryan & Sheu 2016).

from the GLM-based theory to an empirical application on airlines. We then describe our airline data, present our econometric model, and discuss the empirical findings.

4.1 Background on aviation and climate change policy

Airlines currently produce around 2.5% of global CO_2 emissions (McCollum, Gould & Green 2009); as these emissions occur high up in the atmosphere, an effect known as climatic forcing means that the proportion of "effective" emissions is around twice as high. Airline emissions are projected to grow well into the 21st century due to rising global demand for air travel and limited scope for large-scale substitution away from jet engines to low-carbon technologies. As other sectors of the economy, such as electricity generation, decarbonize more quickly the role of aviation in future climate policy is set to grow. Economic regulation appears increasingly likely as countries seek to implement internationally-agreed climate targets in a cost-effective manner.[27]

In this paper, we study the domestic US airline market. This is the world's largest aviation market, producing around 28% of global aviation emissions, but has so far not been subject to carbon pricing. At a baseline carbon price of $50/$tCO_2$, its 2014 emissions of 172 million tCO_2 would have had a value of $8.6 billion. We study the domestic market because international aviation is regulated under a separate organization and set of agreements.

4.2 From economic theory to empirical application

The US aviation market has been largely dominated by several legacy carriers (Alaskan, American, Delta, Hawaiian, United and US Airways) and a large low-cost carrier, Southwest Airlines. Legacy airlines were established before deregulation in 1978; they tend to operate hub-and-spoke networks with relatively high costs and high levels of service while low-cost airlines tend to offer direct flights at lower prices (Borenstein 1992).[28] It is well-established in the literature that market power is an important feature of the deregulated airline industry.

Our unit of analysis is a product, taken to be a route offered by a particular firm (a "carrier-route"). There are important carrier-route heterogeneities that any model should seek to accommodate. First, each carrier-route is its own differentiated product:

[27] Aviation is already subject to carbon pricing in some international jurisdictions. Intra-EU flights have since 2012 been included in the EU ETS. However the impact on airlines has been limited by a low carbon price, mostly in the range of €5-10/tCO_2. Aviation has recently also been included in some regional trading schemes in China, though again with very low carbon prices. As of April 2018, Sweden has introduced an additional carbon tax, on all flights departing from Sweden, that operates alongside the EU ETS. The first global aviation emissions reduction agreement was negotiated by the UN's International Civil Aviation Organization (ICAO) and signed by its 191 member nations in October 2016; it amounts to a carbon-offset scheme for emissions growth after 2020.

[28] A third type of carrier in the US market, the regional airline, is not included in our analysis.

a flight along a given route on Southwest is not the same product as a flight on American (they may differ in leg room, service, complimentary refreshments, air miles, etc.). Hence each carrier-route has a different demand profile (e.g., Berry 1992, Berry & Jia 2010). Second, costs are heterogeneous both across carriers and within a given carrier across its routes: the portfolio of aircraft and the proportion of seats filled vary from carrier-route to carrier-route, with significant impacts on costs (e.g., Berry & Jia 2010, Kwan & Rutherford 2015). Third, the competitive environment varies: the set of competitor firms on each carrier-route is heterogeneous, both in terms of the number of firms and the technologies they use; moreover, the conduct of rivals can vary: one carrier may behave highly competitively on one route and less competitively on another, even facing the same competitors on each route (e.g., Brander & Zhang 1990). Because of these complexities, it is difficult to know with confidence which specific model of competition is best suited to each individual route.

Our objective is to quantify the impacts of carbon pricing in this competitive setting. Leveraging the GLM, we can estimate profit impacts while remaining agnostic about the precise mode of competition across routes. Proposition 1 shows how cost pass-through is a sufficient statistic at the carrier-route level. However, as carbon emissions have been unpriced in US aviation ($\underline{\tau}= 0$), we cannot estimate carbon pass-through rates directly from past data. Instead we estimate fuel cost pass-through rates, and use these to predict carbon cost pass-through. From the viewpoint of a cost-minimizing airline, these costs are equivalent: paying an extra \$1 per unit fuel due to a carbon price is the same as paying an extra \$1 on the price of the fuel. The conversion is simple, since 1 gallon of jet fuel produces 0.00957 tons of CO_2 when burned. We set $\overline{\tau} = \$50/tCO_2$, roughly in line with near-term estimates of the global social cost of carbon (Nordhaus 2017). Importantly, fuel prices have varied substantially over the period we study (2002–2014): the maximum is 540% larger than the minimum; this variation exceeds in magnitude that of a $\$50/tCO_2$ carbon price (given the carbon intensity of jet fuel). Hence our simulated carbon-price shock lies within the range of fuel cost shocks that airlines responded to over the sample period, conferring some external validity on our estimates.[29]

Our empirical approach is based on the assumptions A1–A4 underlying the GLM being appropriate in the context of airline competition. We therefore now discuss these four assumptions in turn:

Input price-taking (A1) is an appropriate assumption in that an airline cannot influence the global oil price, which is the primary determinant of its jet fuel price. Likewise, the price-taking assumption is appropriate in the context of an emissions tax and also

[29]A similar approach of using variation in other input costs to estimate the impact of future environmental costs is also taken by Miller, Osborne & Sheu (2017). In related work, Ganapati, Shapiro & Walker (2016) estimate the pass-through of energy input prices across six US manufacturing industries while Bushnell & Humber (2017) focus on the pass-through of natural gas prices in the fertilizer industry and its implications for the allocation of carbon emissions permits.

approximately correct in the context of a cap-and-trade scheme in which firm i is one among many as in the EU ETS and other current systems.

Cost-minimizing inputs (A2) also appears reasonable for airlines. Fuel costs are often an airline's largest cost, amounting to 20–50% of its total cost base (Zhang & Zhang 2017) so it clearly has strong incentives to minimize fuel costs. Future carbon costs are likely to be managed in conjunction with an airline's overall commodity-market exposure, so we expect these to be similarly optimized. Examples of fuel/emissions reductions by airlines include adjusting flight time, cabin weight, and leasing newer aircraft. These kinds of reversible, continuous and often operational changes are consistent with our framework; anything that airlines did in the past in response to fuel prices, they are likely to do again in response to a carbon price.[30]

Constant returns to scale (A3) is a more substantive assumption, though it is standard in much of the airlines literature.[31] The evidence on whether it holds empirically is inconclusive: while some studies estimate modest scale economies others find no such evidence (Zhang & Zhang 2017) so our analysis is consistent with the notion that these are relatively weak in comparison with the marginal price-cost shifts studied here.[32] Note also that the presence of fixed costs is not an issue for the application of the GLM.

Linear product market behaviour (A4) is the core assumption underlying the GLM. We allow the proportionality factor ψ_i to vary arbitrarily over routes and across carriers. So, in contrast to standard oligopoly models, a higher market share does not necessarily mean a higher profit margin. In this sense, the GLM is quite flexible—and hence also more difficult to reject.

4.3 Description of the data

Our dataset is a panel of price and cost data for airlines over the period 2002Q1–2014Q4.[33] For each carrier i, route j and quarter t, we have the average ticket price p_{ijt}, the average per-person fuel cost k_{ijt}, and a vector of covariates. Hence we have a balanced panel, consisting of $N = 615$ carrier-route observations for $T = 52$ quarters—with a total number

[30] Other abatement activities that fit less well with our approach are one-off, predominantly capital changes, such as purchasing new aircraft or installing wing tips. If these kinds of abatement dominate over the period we study, our historic average pass-through results may give less good predictions of the impact of future carbon pricing.

[31] Brander & Zhang (1990) discuss how to conceptualise constant marginal costs in the case of airlines; Berry & Jia (2010) estimate marginal costs which are constant for a given vector of route characteristics.

[32] There is stronger evidence for economies of *scope*: a higher network density of its route portfolio confers a competitive advantage on an airline—but this is not inconsistent with the GLM theory. As is common in the airlines literature, our empirical analysis considers each route separately, without accounting for potentially complex network effects with other routes (see, e.g., Ciliberto & Tamer 2009). A strength of our estimation procedure is that it allows for arbitrarily complex networks, so long as these are stable over the period we study.

[33] We choose 2002 as the start year to exclude 11 September 2001; our end date avoids the 2015 "mega merger" between American Airlines and US Airways. Appendix D provides robustness checks for different start-end dates.

of observations in excess of 30,000. All monetary quantities are in real 2014Q4 USD.

We construct our data by combining elements of three datasets from the US Bureau of Transportation Statistics. Price data come from the DB1A Origin and Destination Survey, a 10% sample of all airline tickets sold.[34] Prices are for a one-way trip; all round-trip tickets are split equally into two one-way observations. A route is defined by its origin and destination airports, regardless of direction. Our sample is formed by taking all carrier-routes that meet two criteria. First, we examine only direct flights, as is standard in much of the airlines literature.[35] Second, we take only carrier-routes that were continuously operated for the time period we study.[36] This results in a sample containing 36% of the relevant carriers' revenue. We then calculate p_{ijt}, an average of all fares purchased on carrier-route ij in quarter t.

The remaining datasets (T-100 and Form 41) are used to construct k_{ijt}, the average per-passenger fuel cost for flying with carrier i on route j in time t.[37] We construct this by averaging carrier-route ij's aircraft-specific fuel expenditure, weighted by the proportion of flights operated by that aircraft type on carrier-route ij, and then divide by the proportion of filled seats. (Appendix B describes this procedure in more detail.) The value of k_{ijt} is determined by three factors: (i) the market price of jet fuel, which tracks the crude oil price; (ii) the fuel efficiency of the passenger's journey, as driven by the type and age of aircraft used, the configuration the seating and the proportion of seats filled (any other variation in the airline's physical operating procedures can also influence this factor); and (iii) the carrier's use of hedging or other financial products when buying fuel. This varies significantly between carriers and over time for a given carrier. For example, in our sample period, Southwest was known for its extensive use of hedging, while US Airways never hedged. Carriers therefore ended up paying very different prices: in 2008 (when oil prices were rising) US Airways paid 30% more for each gallon of fuel than Southwest, whereas in 2009 (when oil prices fell) it paid 18% less.[38]

[34]We use a cleaned version of the DB1A provided by Severin Borenstein. The following ticket types are excluded: international, first class, frequent flier (those with a price less than $20), entry errors (price higher than $9,998 or five times the industry standard for that route-time), and open or circular itineraries. Observations are aggregated up to the carrier-route-time level. (In Appendix A.2, we provide a microfoundation for the widespread approach in the literature of analyzing economy-class tickets in a separable way from business or first-class tickets.)

[35]Indirect flights—involving a change of aircraft at another airport en route, using the airline's hub-and-spoke network—are well-known to have different economic characteristics to direct flights. Excluding indirect flights is, therefore, commonplace in the airlines literature (e.g., Borenstein & Rose 1994; Goolsbee & Syverson 2008; Gerardi & Shapiro 2009).

[36]We consider a route to be operated by a carrier in any quarter where it carried at least 1,000 passengers. This is equivalent to one small (83 passenger) service per week. (Because of the lag structure in our specification, this requires the route to be continuously operated over the slightly longer period 2000Q1-2014Q4.)

[37]The overlap between the DB1A and T-100 is good but not perfect (see Goolsbee & Syverson 2008 for a fuller discussion). Merging with data from T-100 results in around 10% of DB1A revenue being dropped.

[38]We do not have detailed data on the precise extent of hedging by each carrier at each point in time. Turner & Lim (2015) document and analyse the different hedging strategies of US airlines, and the

	Southwest				Legacy			
	mean	s.d.	min	max	mean	s.d.	min	max
Price ($)	157.31	40.52	74.78	298.91	230.82	78.21	52.14	683.50
Fuel cost ($)	29.22	15.69	5.29	101.52	50.08	31.05	2.33	366.63
Distance (miles)	688	407	148	2,106	1,097	706	84	3,784
Emissions (tCO_2)	0.13	0.06	0.03	0.44	0.21	0.11	0.02	1.18
Emissions cost ($)	6.70	2.92	1.71	21.92	10.47	5.54	1.12	59.12
Passengers (000s)	195	172	5	1,172	153	135	4	1,263
No. firms	3.28	2.41	1.00	17.00	3.67	2.24	1.00	17.00
Fraction seats filled	0.72	0.10	0.33	0.97	0.79	0.10	0.23	0.97
Revenue ($ million)	24.76	18.78	0.83	135.07	28.99	24.92	0.33	238.11
Revenue in sample	0.42	–	–	–	0.34	–	–	–
No. routes	212	–	–	–	403	–	–	–
No. observations	11,024	–	–	–	20,956	–	–	–

Table 1: Average ticket prices, per-passenger fuel cost and other carrier characteristics.

Notes: All results are for the period 2002Q1-2014Q4. Price, fuel cost, emissions and emissions cost are per passenger. Passenger numbers and revenue are year averages. Emissions cost are calculated at a carbon price of $50/$tCO_2$.

Having constructed p_{ijt}, k_{ijt} and a set of covariates (described below), we keep data for the six largest legacy carriers and for Southwest.[39] The routes in our sample make up 25% by revenue of all domestic US aviation activity over the period.[40]

Table 1 presents descriptive statistics on airlines' prices, costs and other variables related to competition and environmental performance. Given that their characteristics are quite similar, but distinct from Southwest, we combine the six legacy carriers into a single group; see Table 9 in Appendix C for the breakdown by legacy carrier. Southwest tends to fly shorter routes than the legacy carriers; it charges lower prices and has lower fuel costs and emissions. Southwest is comparable to the average legacy carrier in terms of size, capacity utilization, and the number of competitors it faces on its routes.

Figure 1 shows trends over the period across all carriers. Panel (a) compares the average per-passenger fuel cost k_t with the spot price of jet fuel. They track each other

different effective fuel prices that result.

[39] This is the result of dropping three very small carriers: one regional (ExpressJet), and two low cost carriers (Frontier and Spirit). These airlines have an average of only 12 routes each.

[40] The 7 large carriers have a market share by revenue of 61% over the period (this rises to 69% at the end of the period due to mergers). The proportion of each carrier's revenue included in our sample is shown in Table 1, with the heterogeneity coming principally from differences in the use of direct vs hub-and-spoke business models.

very closely; the slight lag of fuel costs behind the spot price is principally due to hedging. The figure also reveals substantial variation in fuel costs over the period. Panel (b) plots average ticket prices (left axis) against per-passenger fuel (right axis). As expected, there is a positive correlation between price and fuel cost.

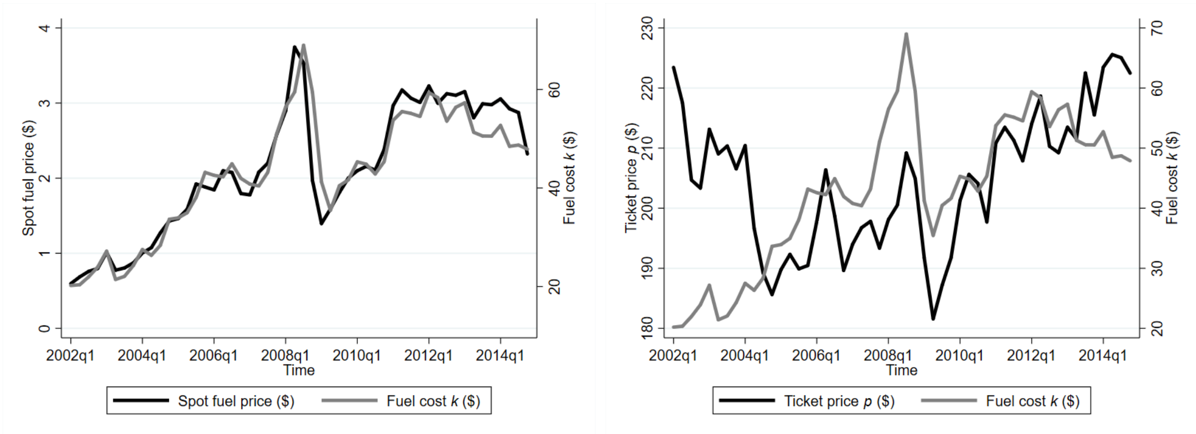

(a) (b)

Figure 1: (a) Spot price of jet fuel and per-passenger fuel cost; (b) Ticket price and per-passenger fuel cost.

Notes: Variables are quarterly averages over all carrier-routes in our sample.

4.4 Baseline econometric specification

Following a standard approach in the pass-through literature (e.g., Fabra & Reguant 2014, Stolper 2016, Atkin & Donaldson 2015, and Miller, Osborne & Sheu 2017), we regress prices on costs and control variables, giving the following baseline specification:

$$p_{ijt} = \sum_{m=0}^{3} \rho_{ij}^{m} k_{ij,t-m} + X'_{ijt}\beta_{ij} + \epsilon_{ijt}, \tag{7}$$

where we are interested in the "equilibrium" rate of fuel cost pass-through for carrier-route ij, $\rho_{ij} = \sum_{m=0}^{3} \rho_{ij}^{m}$. We include the current per-passenger fuel cost k_{ijt} with 3 quarters of lags, thus allowing for fuel cost adjustment to take up to one year. Hence pass-through ρ_{ij} measures the price increase one year after a permanent \$1 increase in fuel costs.[41]

Equation (7) allows the pass-through rate to vary across carriers and routes, but is fixed over time. In this respect and others it is similar to the approach in Atkin & Donaldson (2015), who also give a useful discussion of the implications of assuming a pass-through rate is constant over time.

[41] Our choice of specification reflects the fact that prices can be quickly and fully adjusted by airlines. Hence we do not include a lagged dependent variable, reflecting a "speed of adjustment", such as in many dynamic econometric specifications.

Our specification does not directly include competitors' costs. Clearly, however, a firm's equilibrium price is usually a function of its own as well as its rivals' costs. The literature on pass-through estimation takes a variety of approaches to this. Miller, Osborne & Sheu (2017) include in their overall pass-through rate the coefficient on a measure of firm-specific competitive pressure to capture the rival-cost component of pass through. Stolper (2016) argues for the theoretical importance of including rivals' costs but then, for reasons of identification, estimates an overall industry cost pass-through. Our approach is to allow firm i on route j to have costs k_{ijt} that may be correlated in any arbitrary way with any other rival l's costs. The substantive assumption is that, however own and rivals' costs are related, this is unchanging over the period we study. On this basis, for carrier-route ij, the overall pass-through rate ρ_{ij} we obtain contains both own-cost and rival-cost effects.[42]

Carrier-specific pass-through rates ρ_i are obtained from specification (7) by running a separate regression for each carrier-route, and then taking a weighted average of the ρ_{ij} to obtain a pass-through rate at the carrier level, ρ_i. (The weights are explained below.) This imposes no homogeneity restrictions on the parameters across carrier-routes, which is important given the heterogeneities discussed above. In running a separate regression for each product, we take a similar approach to Atkin & Donaldson (2015). The procedure could also be considered a special (non-dynamic) case of the "Mean Group" estimator in Pesaran & Smith (1995). A fixed effects specification (which we report for completeness in Appendix D), by contrast, would impose homogeneity on all parameters other than the constant.[43] Note that allowing pass-through rates and other parameters to vary across carrier-routes does not mean the routes are independent in an economic sense, rather that their interdependencies are one of the many characteristics captured by the pass-through rate that we seek to estimate.

Equation (7) cannot be estimated using OLS because the per-passenger fuel cost k_{ijt} is potentially endogenous. It depends on the number of passengers flying in quarter t, which in turn will generally be an outcome of the price p_{ijt}. So our dependent variable includes in its denominator an outcome variable, introducing the possibility of simultaneity bias. To address this, we use the spot price of jet fuel as an instrument for fuel cost k_{ijt}. Since the price of jet fuel is determined by the global oil price, it is exogenous to passenger numbers on a particular route, satisfying the exclusion principle. In order to accommodate potential hedging, we use 8 quarters of spot prices to estimate the 4 quarters of fuel costs

[42] In this respect our approach is similar to both Atkin & Donaldson (2015) and Stolper (2016). Fabra & Reguant (2014) take another apporach, given their focus on electricity: they analyze pass-through for a homogeneous good (electricity) with a single market price, and their resulting pass-through rate gives the dollar increase in the electricity price when the carbon costs of the marginal supplier increase by $1.

[43] Estimating equation (7) as a system of N seemingly unrelated regression equations would not change the point estimates but would give efficiency gains. SURE estimation is, however, not feasible here because the number of equations in our system is so much larger than the degrees of freedom of each individual regression.

in equation (7). Hence, the first stage regressions are given by:

$$k_{ij,t-m} = \sum_{q=0}^{7} \gamma_{ij}^{m,q} f_{t-q} + X'_{ijt}\beta_{ij}^{m} + \epsilon_{ijt}^{m} \quad \text{for each } m \in \{0,1,2,3\} \tag{8}$$

where f_t is the spot fuel price.[44]

Equations (7) and (8) also includes a vector of controls X_{ijt} which capture changes in supply and demand that could otherwise introduce omitted variable bias into the regression. In addition to a constant, X_{ijt} is made up of the following: (i) We include GDP growth to proxy for demand because jet fuel prices closely track the oil price, which may be systematically related to demand for air travel. We construct a route-specific measure by taking the average GDP of the two states at either end of the route. (ii) We construct an index of certain key (non-fuel) costs at the carrier level principally made up of labour and aircraft maintenance costs. (iii) We include the number of competitor firms on route j. If this quantity remained stable over time, there would be no need to include it in regression (7); however we do see entry and exit in the data, which could potentially be systematically related to fuel costs. (iv) We also include the number of potential entrants. This is defined as the number of airlines operating services out of both the origin and destination of route j, but not actually connecting these with a direct flight. Goolsbee & Syverson (2008) show that potential entrants have a significant effect on incumbents' pricing. (v) Finally, we include quarterly dummies to control for any seasonality in pricing. We report Newey-West standard errors throughout, which are heteroskedasticity and autocorrelation consistent.

4.5 Illustration for an individual carrier-route

Before presenting the full results, we illustrate our approach to estimating pass-through rates for a single route, Southwest's service between Phoenix, Arazona (Phoenix Sky Harbor International Airport, PHX) and San Antonio, Texas (San Antonio International Airport, SAT). Table 2 (a) gives average prices and costs for this route; the comparison with Table 1 shows that it is reasonably representative of Southwest's portfolio of routes in our dataset.[45] The average ticket price is \$200, and the average fuel cost is \$33 per passenger. Figure 2(a) shows how the ticket price p_{ijt} changed with the fuel cost k_{ijt} over time. The causal component of this relationship is what we seek to estimate. Figure 2(b) shows the movement of the covariates in X_{ijt} over the period.

Estimating equation (7) for i = Southwest and j = PHX-SAT with 2SLS, using first stage regression equation (8), gives an equilibrium rate of pass-through of $\rho = 1.38$ ($\pm$

[44] We use jet fuel price data from Bloomberg: JETINYPR index (New York Harbor 54-Grade Jet Fuel).

[45] The route a little longer than average for Southwest; it is correspondingly a little more expensive in terms of terms of ticket price and costs. No other carriers in our sample operate this route continuously over the period 2002-2014.

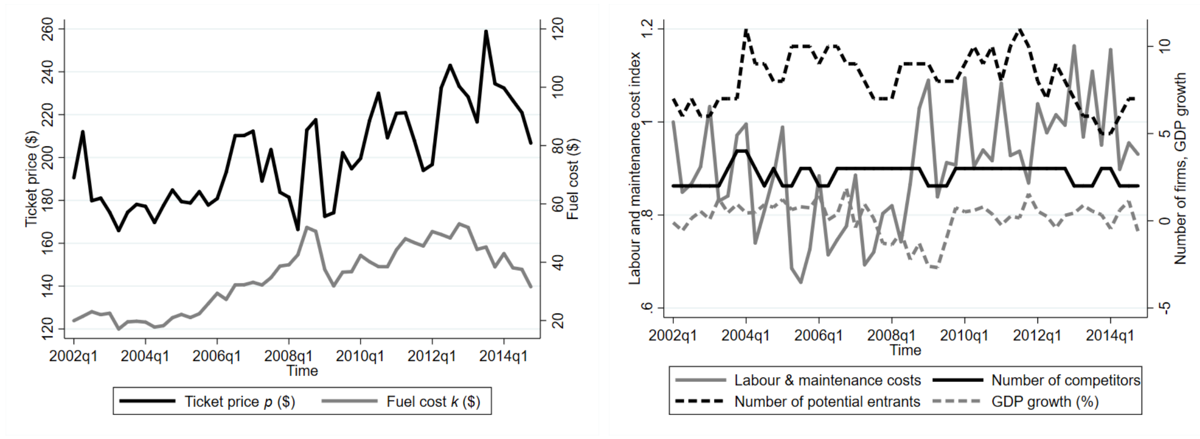

(a) (b)

Figure 2: (a) Ticket price and per-passenger fuel cost for Southwest on the route from Phoenix (PHX) to San Antonio (SAT); (b) Southwest covariates on the same route.

Notes: Covariates are described in Section 4.4.

0.62). (Throughout the text, 95% confidence intervals are given in parentheses.) Table 7 in Appendix C.1 contains the full regression results. The covariates generally have their expected signs, although some are not significantly different from zero; the relatively large standard errors are not surprising for a regression with 52 observations and 19 parameters to estimate. Looking ahead to the full results, the standard errors will fall dramatically as we bring in many more observations.

Up to this point, we have made no particular assumptions about Southwest's behaviour on this route, other than the standard assumptions inherent in running a linear regression. Now invoking the GLM's A1–A4 and using Proposition 1(a), we conclude that Southwest's pass-through in excess of 100% implies its sales and profits rise with the carbon price; in equilibrium, this implies that it gains market share from rivals on this route. Table 2 (b) summarizes the predicted impact of a \$50/tCO_2 carbon price implemented at the end of the period (2014Q4). At that time, emissions on this route were 0.13 tCO_2 per passenger per flight so this translates into a carbon cost shock of \$6.40 per passenger. Given the pass-through rate $\rho = 1.38$, Proposition 1 shows that Southwest's profits rise by \$0.38 ($\pm$ 0.62) million per year, or +2.22% ($\pm$ 3.59) of revenue on this route.

4.6 Estimation results for all routes

We now present estimates of pass-through rates for all routes averaged at the carrier level. We estimate specification (7) by 2SLS for each carrier-route, to obtain pass-through rates ρ_{ij}. Using Proposition 1(a), the profit impact of a carbon price τ on carrier i across all its j routes is given by $\Delta\Pi_i \simeq \tau \sum_j 2e_{ij}(0)(1-\rho_{ij}) = 2(1-\rho_i)\tau e_i(0)$, where ρ_i is its

(a) Descriptive statistics	
Price ($)	200.32
Fuel cost ($)	32.59
Number of firms	2.57
Number of potential entrants	8.10
Distance (miles)	843
Passengers, annual	76,014
Proportion of seats filled	0.73
Revenue in 2014 ($ million)	17.36
Emissions in 2014 (tCO_2)	0.13
Emissions cost in 2014 ($)	6.40

(b) Regression results	
Pass through	1.38
	(0.32)
Profit impact (% of revenue)	2.22
	(1.83)
R squared	0.68
No. observations	52

Table 2: (a) Descriptive statistics, and (b) Pass-through estimates for Southwest's route from Phoenix (PHX) to San Antonio (SAT).

Notes: Descriptive statistics are averages over the 2002Q1-2014Q4 sample period. Emissions cost calculated using $50/$tCO_2$ carbon price. Results for 2SLS estimation of (7) for PHX-SAT. Profit impact based on $50/$tCO_2$ carbon price implemented at 2014Q4 using 2014 year average emissions per passenger.

emissions-weighted average pass-through rate.[46] We use 2014 emissions to compute the weights, as we want to predict the profit impact of a carbon price introduced at the end of the period.

Table 3 summarizes the results. The industry average pass-through rate is 0.78 ($\pm$ 0.10) lies below 100%. But we find strong evidence of intra-industry pass-through heterogeneity: Southwest has a weighted average cost pass-through of $\rho = 1.48$ ($\pm$ 0.08), compared to the legacy carrier average of $\rho = 0.55$ ($\pm$ 0.12).[47, 48] Table 8 in Appendix C.2 gives more detailed regression results, and shows that every covariate has the expected sign, and most are highly significant. Figure 3 plots the distributions of pass-through estimates for each group.[49] Assuming that the GLM's assumptions A1-A4 hold for each carrier-route, we estimate that a $50/$tCO_2$ carbon price raises Southwest's profits by

[46] As described in Footnote 32, our approach stays agnostic on the presence ot network effects, as these are one of the things that may impact pass-through rates. The econometric approach followed here does, however, assume the networks are stable over time.

[47] The carrier-level standard errors are calculated on the assumption that each pass-through rate is estimated with error, but that these errors are independent of each other. Hence the variance of the weighted average $\rho_i = \sum_{ij} \alpha_{ij}\rho_{ij}$, reduces to $\mathrm{Var}[\rho_i] = \sum_j \alpha_{ij}^2 \mathrm{Var}[\rho_{ij}]$, where $\alpha_{ij} = \frac{e_{ij}}{\sum_j e_{ij}}$. In so far as any correlation between the ρ_{ij} for an airline i is likely to be positive, the standard errors calculated under the assumption of no correlation are a lower bound on the true standard error. As noted in Footnote 43, SURE is infeasible so we cannot obtain the covariance matrix for the pass-through estimates.

[48] Table 10 in Appendix C.3 gives results for individual legacy airlines. There is substantial heterogeneity also amongst them, though all have firm-level pass-through rate below 100%.

[49] The distribution of pass-throughs does not necessarily have a clear economic interpretation, as each is measured with (a different) error. For example, it is unlikely an outlier in the legacy distribution is a carrier-route that really has $\rho = -5$. Atkin & Donaldson (2015) also find a distribution of pass-through rates with a 'sensible' mean, but some extreme (for example negative) outliers. See Pesaran, Shin & Smith (1999) for a discussion of this issue from the econometrics literature.

+2.95% (± 0.44) of revenue. In stark contrast, the legacy carriers would suffer an average profit impact of –3.56% (± 1.02) of revenue.[50]

	Southwest	Legacy	All
Pass through	1.48	0.55	0.78
	(0.04)	(0.06)	(0.05)
Profit impact (% revenue)	2.95	-3.56	-1.59
	(0.22)	(0.51)	(0.36)
Profit neutral permit allocation	-0.96	0.90	0.43
	(0.07)	(0.13)	(0.10)
R squared	0.67	0.45	0.56
No. routes	212	403	615
No. observations	11,024	20,956	31,980

Table 3: Pass-through estimates and predicted profit impacts across carrier types.

Notes: Results for 2SLS estimation of (7) for all carrier-routes, averaged using 2014 carrier-route emissions as the weights. Profit impacts based on \$50/tCO$_2$ carbon price implemented in 2014Q4.

Making the admittedly expedient assumption that the pass-through rates obtained from our sample of continuously operated routes are representative of *all* routes operated by these airlines, we cautiously calculate a total profit impact of +\$0.51 (± 0.07) billion for Southwest and –\$1.46 (± 0.41) billion for the legacy carriers. [51] This compares with the companies' reported 5-year-average annual profits of, respectively, \$1.17 billion and \$4.26 billion in 2014. Hence our results suggest a \$50/tCO$_2$ carbon price would increase Southwest's profits by +44% while cutting those of the legacy carriers by 34%.

To the best of our knowledge, this finding of pass-through heteroegeneity in the airline industry is novel. It is, however, consistent with earlier findings in the literature. Goolsbee & Syverson (2008), Ciliberto & Tamer (2009), and Berry & Jia (2010) variously point to Southwest being more efficient, better able to cope with shocks, or especially threatening to its rivals. Gaudenzi & Bucciol (2016) report jet fuel price rises are associated with significantly more negative stock-market returns for legacy carriers than for Southwest.[52]

[50]It is worth noting that, despite its higher pass-through rate, Southwest's price remains below that of the legacy carriers. The weighted average emission cost shocks are \$6.48 and \$7.57 for Southwest and the Legacy carriers respectively, given their 2014 emissions. This translates in to price rises of \$9.74 and \$6.41, respectively, to give final prices of \$187.03 and \$249.68. Note that this simple comparison of means is indicative only, as the portfolio of routes is different for each airline (see below for further discussion).

[51]It is unlikely that the continuously operated routes have identical pass-through rates to routes closed or newly opened over the period 2002-2014. The average would be similar if under- and over-estimates coming from closed and opened routes (or vice versa) approximately cancelled each other out. Similarly there is no reason to believe direct and indirect routes would have identical pass-through rates. Hence, the extrapolated results should treated as indicative only.

[52]Gaudenzi & Bucciol (2016) stress differences in hedging strategies; our empirical findings here, in

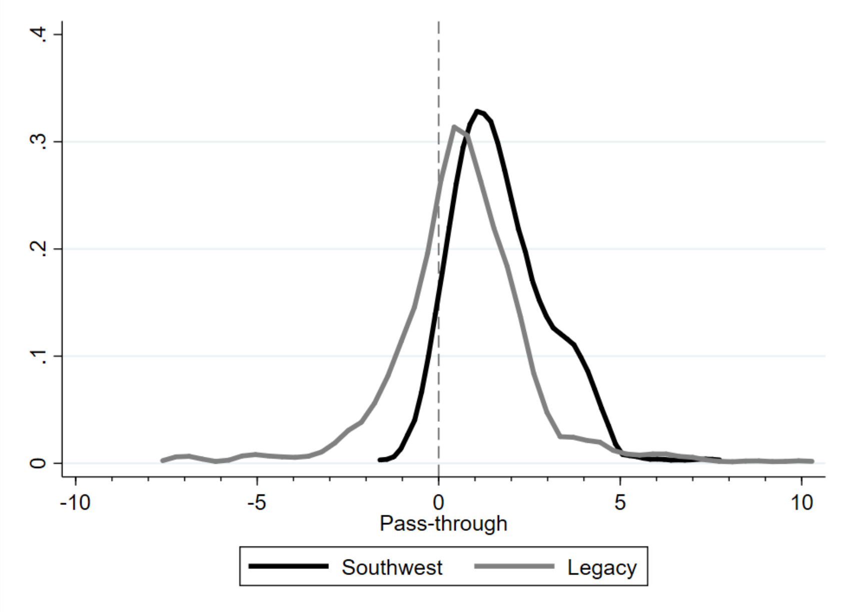

Figure 3: Kernel density of pass-through estimates across routes for Southwest and legacy carriers.

The final result summarized in Table 3 is the profit-neutral permit allocation (Hepburn, Quah & Ritz 2013). This is the proportion of a firm's pre-regulation level of emissions that would need to be freely allocated to preserve its pre-regulation profits. We find that, due to their low pass-through, legacy carriers would need 90% (±26%)—that is, almost all—of their permits for free. By contrast, Southwest's profit-neutral allocation is strongly negative—which is the flipside of its profits rising with carbon regulation. Our industry-wide profit-neutral allocation of 43% is considerably higher than previous empirical findings for sectors such as electricity and cement, suggesting that aviation gets hit relatively hard by carbon pricing.

Appendix D details robustness checks and extensions to the baseline empirical results.

4.7 What explains the heterogeneity in pass-through rates?

Table 1 revealed that Southwest tends to operate routes that are shorter than the legacy carriers. Long haul flights are likely to have systematically different characteristics: on the demand side, there are fewer close substitutes (such as bus or rail travel); on the supply side, entry may be more difficult so there may be greater market power (e.g., Brander & Zhang 1990; Berry & Jia 2010). We next present a decomposition of the difference in pass-through rates based on three factors: route portfolio, production costs, and product differentiation.

First, to quantify how route characteristics drive the overall difference in airlines' pass-through rates, we recalculate carrier-level averages for the subset of routes that serve a

effect, offer a new explanation for their results.

market which is common to both airline types. We compute an average pass-through that is not weighted by emissions, so the importance of a route in an airline's overall portfolio is not impacting this measure. The results for the 49 common routes are reported in Table 4; the unweighted average pass-throughs are somewhat higher than our previous weighted average results. Moving to common routes narrows the difference in pass-through rates considerably, mainly because legacy pass-through is much higher. We can attribute 61% of the original difference in pass-through rates in Table 3 to differences in the airlines' route portfolios.[53]

	Southwest			Legacy		
	All weighted	All un-weighted	Common un-weighted	All weighted	All un-weighted	Common un-weighted
Pass through	1.48	1.72	1.61	0.55	0.69	0.98
	(0.04)	(0.04)	(0.09)	(0.06)	(0.06)	(0.18)
No. routes	212	212	49	403	403	49

Table 4: Comparison between Southwest and legacy carriers of pass-through estimates for common routes (operated by both Southwest and a legacy carriers) with the full sample of routes.

Notes: The move from pass-through weighted by the emissions of each carrier-route to unweighted pass-through yields a like-for-like comparison across carriers.

To further explore these differences, we present in Table 5 a cross-tabulation of weighted average pass-through rates. Taking the full sample of carrier-routes, we calculate the triciles (i.e., separated at the 33rd and 67th percentiles) of route distance and number of firms (n) operating the route. This cross-tabulation shows that Southwest has significantly higher pass-through rates on its shorter routes, however competitive they are (as proxied by n).[54] Indeed, the number of firms has no clear relationship with pass-through rate for either carrier type.[55] This suggests that it is differences in demand patterns associated with flight distance between Southwest and legacy carriers—more than differences in supply conditions—that are driving pass-through asymmetries.[56]

Second, we examine the effect of production cost asymmetry. Southwest has lower average fuel costs on like-for-like routes: \$30.14 per passenger compared to \$36.02 for the

[53]The proportion of the difference in pass-through on the common (unweighted) routes (0.63) is 61% of the difference in pass-through on all (unweighted) routes (1.03).

[54]The relationship between distance and both demand and supply has been studied for a long time in the literature (Brander & Zhang 1990; Berry & Jia (2010).

[55]A similar conclusion comes out of running a regression of estimated pass-through rates on distance and the number of firms: the coefficient on distance is significantly negative, whereas the coefficient on n is not statistically different from 0. Using the Herfindahl index instead of the number of firms as a measure of competition yields similar results from both the cross-tabulation and regression.

[56]A limitation of our GLM-based approach is that we are not able to identify the underlying causal factors driving the results to the extent that a full structural model may be able to.

	Southwest			Legacy		
	Short *distance* $\in$ $[0, 570)$	Medium *distance* $\in$ $[570, 1034)$	Long *distance* $\in$ $[1034, 3784]$	Short *distance* $\in$ $[0, 570)$	Medium *distance* $\in$ $[570, 1034)$	Long *distance* $\in$ $[1034, 3784]$
Small $n \in [1, 2.3)$	2.00 (0.10) *34*	1.03 (0.07) *30*	0.80 (0.07) *24*	1.03 (0.22) *39*	0.26 (0.29) *29*	0.73 (0.09) *49*
Medium $n \in [2.3, 4.3)$	2.48 (0.10) *35*	0.90 (0.09) *19*	0.60 (0.08) *11*	0.58 (0.31) *34*	0.78 (0.21) *56*	0.00 (0.12) *53*
Large $n \in [4.3, 12.5]$	2.55 (0.10) *33*	0.87 (0.09) *20*	0.64 (0.16) *7*	-0.18 (1.28) *27*	0.87 (0.12) *60*	0.68 (0.08) *59*
All n	2.40 (0.56) *102*	0.91 (0.38) *68*	0.70 (0.33) *42*	0.46 (2.35) *99*	0.75 (1.14) *143*	0.46 (0.59) *161*

Table 5: Cross-tabulation of pass-through estimates by route distance and number of competitors.

Notes: The columns are triciles of route distance (in the full sample). The rows are triciles of the average number of firms operating a route (in the full sample). Standard errors in parentheses; number of routes in italics.

legacy carriers. This is principally due to use of newer aircraft and more efficient seating configurations. In the absence of any demand asymmetries between firms there is a single market price; in such cases, any two firms' (i and j) pass-through rates are related only by their relative cost shocks, $\rho_j/\rho_i = z_i(\tau)/z_j(\tau)$. Southwest's superior fuel efficiency can thus explain 68% of the remaining difference in pass-through rates on common routes, or 26% of the original difference in Table 3.[57]

Third, 32% of the pass-through differential on common routes remains to be explained (corresponding to 12% of the original difference). This residual difference now shows heterogeneity in pass-through rates for a uniform cost shock on common routes (that hits all carriers equally). We therefore attribute this residual to demand-side asymmetries between carriers, based on their differentiated-product offering. Such demand asymmetries now imply departures from a single market price on common routes so also mean that differences in competitive conduct may be driving pass-through asymmetry.[58]

[57] All else equal, we would expect the legacy carriers to have a pass-through rate 16% lower than Southwest using $\rho_j/\rho_i = z_i(\tau)/z_j(\tau)$. Using Southwest as the baseline, this explains 26% of the original difference in Table 3. (Alternatively, legacy pass-through could be used as the baseline, which would yield an answer of 12%.)

[58] For example, Southwest has a higher market share on the common routes in our sample; this is one factor than can lead to asymmetry in pass-through in the context of differentiated-products competition.

In sum, this decompostion suggests that around 60% of the pass-through difference is due to different route portfolios; around two thirds of the remaining difference is driven by Southwest's superior fuel efficiency and the final third stems from differences in demand patterns (and competitive conduct).[59]

4.8 Further results on the determinants of pass-through

To complete our empirical analysis, we next present three additional results on the determinants of pass-through. These are each based on an extension of the baseline specification (7) that includes an interaction term on fuel costs. We here summarize our findings; see Appendix E for details including Table 12 for the numerical results.

First, we interact legacy carriers' fuel costs with a dummy equal to 1 if it faces *actual* competition from Southwest on a route. We find that legacy pass-through is significantly lower when facing direct competition from Southwest. However, when we also include a dummy for *potential* entry of Southwest, then this effect disappears—and is replaced by a strongly negative effect of potential competition. Thus we find new evidence that the threat of entry by Southwest affects not only the level of incumbent pricing (Goolsbee & Syverson 2008) but also their cost pass-through.

Second, we interact airlines' fuel costs with a bankruptcy dummy. Most legacy carriers went through spells of bankruptcy during our period of study, including American, Delta and United (see also Borenstein 2011). There is prior evidence that airlines' financial constraints can lead to price wars (Busse 2002). We find that legacy carrier cost pass-through is 15% higher, on average, during bankruptcy. This new evidence is consistent with the intuition that a firm in severe financial distress is less able or willing to absorb cost shocks.

Third, in the spirit of Kellogg (2014), we interact fuel costs with simple measure of fuel-cost volatility. We find that higher input-cost volatility reduces cost pass-through, both for Southwest and the legacy carriers. This finding appears to be new to the pass-through literature. It also carries a potentially important insight for instrument choice in climate policy. A cap-and-trade scheme creates volatile carbon prices while a carbon tax does not. Hence, our finding suggests that the profit impact, at the same average carbon price, of a cap-and-trade scheme may be more negative than that of a tax.

5 The political equilibrium carbon price

We now present an application to the political economy of regulation. The GLM allows us to bring together two strands of literature: (1) an influential literature following Grossman

[59] To add together the above elements of this back-of-the-envelope decomposition, we implicitly require the ratio of fuel efficiencies $z_i(\tau)/z_j(\tau)$ between Southwest and the lagacy carriers to be the same for the the whole sample as it was in the common routes.

and Helpman (1994) in which firms lobby the government over the magnitude of regulation and its policy is "for sale"; (2) the classic literature following Buchanan (1969) on second-best emissions taxation under imperfect competition.[60] We show that both models are driven, in the same way, by a firm-weighted rate of cost pass-through. We then estimate the "political equilibrium" (third-best) carbon price for US airlines.

5.1 A model of equilibrium regulation with political lobbying

As in Grossman and Helpman (1994), consider a partial-equilibrium setting in which the government cares about social welfare W as well as (aggregate) political contributions by $n \geq 2$ regulated firms. Let $K_i(\tau)$ denote firm i's political contribution to the government as a function of the (common) emissions price τ. The government's payoff is:

$$U_{\text{gov}}(\tau) = W(\tau) + \lambda \sum\nolimits_{i=1}^{n} K_i(\tau),$$

where the parameter $\lambda \geq 0$ measures the government's openness to lobbying. For $\lambda = 0$, the government's problem will boil down to standard welfare-maximization; for larger λ, its regulatory policy is increasingly for sale.

Following Bernheim and Whinston (1986) and Grossman and Helpman (1994), the equilibrium of the lobbying game is for each firm i to offer a contribution function $K_i(\tau) = \Pi_i(\tau) + u_i$, where u_i is a constant. Substituting this into the government's payoff function, the first-order condition for the political equilibrium emissions price $\tau^{\star}(\lambda)$ is given by:

$$\frac{dU_{\text{gov}}(\tau)}{d\tau} = \frac{dW(\tau)}{d\tau} + \lambda \sum\nolimits_{i=1}^{n} \frac{d\Pi_i(\tau)}{d\tau} = 0. \tag{9}$$

We assume that this problem is well-behaved, and focus on the interesting case of an interior solution with $\tau^{\star}(\lambda) > 0$. As will become clear, this includes the usual property that an emissions tax is successful at reducing aggregate emissions, $dE(\tau)/d\tau < 0$.

To make further progress, additional assumptions are needed. First, we assume that A1–A4 from the GLM now hold for *all* firms in the industry. Second, for simplicity, we take each firm's emissions intensity to be fixed (so that $z_i(\tau) = z_i(0)$), while still allowing for arbitrary heterogeneity across firms. Third, we assume that consumers are utility-maximizers, with aggregate consumer surplus $S = V(x_1, ..., x_i, ..., x_n) - \sum_{i=1}^{n} p_i x_i$, where $V(\cdot)$ is gross consumer utility; unlike much of the literature on second-best emissions pricing (Requate 2006), this allows for firms' products to be horizontally and/or vertically differentiated. Fourth, environmental damages depend on aggregate emissions $D(E)$, where $E = \sum_{i=1}^{n} e_i$, and, as usual, are increasing and convex, $D'(\cdot), D''(\cdot) > 0$.

The timing of the game is as follows. First, firms choose their contribution functions

[60] By assumption, the government does not have access to another policy instrument (such as a price control) to directly address market power.

K_i. Second, government sets the emissions price τ. Third, firms compete according to the GLM (now taking τ as given, as per A1).

Social welfare can therefore be written, in terms of the emissions tax, as:

$$W(\tau) = V(x_1(\tau), ..., x_i(\tau), ..., x_n(\tau)) - \sum_{i=1}^{n} C_i(x_i(\tau)) - D(E(\tau)),$$

reflecting that firms' revenues $\sum_{i=1}^{n} p_i x_i$ are a transfer from consumers and firms' emissions costs τE are a transfer to government.

As a benchmark, recall that the standard Pigouvian tax to address the environmental externality under perfect competition (that is, marginal utility equals price by consumer optimization, and price then equals marginal cost) is to set the emissions price equal to the social marginal damage (social cost of carbon), $\tau = D'(E(\tau))$.

We obtain the following characterization of the equilibrium degree of regulation:

Proposition 2 *At an interior solution, the political equilibrium emissions tax satisfies:*

$$\tau^{\star}(\lambda) = \left[\frac{D'(E(\tau))}{1 + \dfrac{(1+2\lambda)}{-\eta(\tau)} \sum_{i=1}^{n} \dfrac{e_i(\tau)}{E(\tau)} [1 - \rho_i(\tau)]} \right]_{\tau = \tau^{\star}(\lambda)},$$

where $\eta \equiv [dE(\tau)/E(\tau)] / [d\tau/\tau] < 0$ *is the tax elasticity of industry emissions. The political equilibrium tax is less than the Pigouvian tax if and if only industry-weighted pass-through is less than 100%,*

$$sign\left\{D' - \tau^{\star}(\lambda)\right\} = sign\left\{\sum_{i=1}^{n} \frac{e_i(\tau)}{E(\tau)} [1 - \rho_i(\tau)]\right\},$$

and, if below the Pigouvian tax, it decreases with the government's openness to firm lobbying, $d\tau^{\star}(\lambda)/d\lambda < 0$.

Proof. We first derive the marginal impact of a tax on welfare, and then use this to pin down the equilibrium tax rate. By consumer optimization, marginal utility equals price, $\partial V/\partial x_i = p_i(\tau)$, and, by A3 from the GLM, $\partial C_i/\partial x_i = k_i(\tau)$. Using this, differentiation of the welfare function yields:

$$\frac{dW(\tau)}{d\tau} = \sum_{i=1}^{n} [p_i(\tau) - k_i(\tau)] \frac{dx_i(\tau)}{d\tau} + (\tau - D'(E(\tau)) \frac{dE(\tau)}{d\tau}.$$

A4 $x_i(\tau) = \psi_i[p_i(\tau) - k_i(\tau)]$ implies that $dx_i(\tau)/d\tau = \psi_i[\rho_i(\tau) - 1] z_i(0)$ since $dk_i(\tau)/d\tau = z_i(0)$ (as firms' emissions intensities are assumed constant, with $z_i(\tau) = z_i(0)$). Since also $e_i(\tau) = z_i(0) x_i(\tau)$, this leads to a welfare impact:

$$\frac{dW(\tau)}{d\tau} = \sum_{i=1}^{n} e_i(\tau)[\rho_i(\tau) - 1] + (\tau - D'(E(\tau)) \frac{dE(\tau)}{d\tau}.$$

The profit impact follows directly from the proof of Proposition 1(a):

$$\sum\nolimits_{i=1}^{n} \frac{d\Pi_i(\tau)}{d\tau} = 2\sum\nolimits_{i=1}^{n} e_i(\tau)[\rho_i(\tau) - 1],$$

Putting these parts together in the first-order condition from (9) shows that $\tau^{\bigstar}(\lambda) > 0$ is determined by:

$$\left.\frac{dU_{\text{gov}}(\tau)}{d\tau}\right|_{\tau=\tau^{\bigstar}(\lambda)} = \left[(1+2\lambda)\sum\nolimits_{i=1}^{n} e_i(\tau)[\rho_i(\tau) - 1] + (\tau - D'(E(\tau))\frac{dE(\tau)}{d\tau}\right]_{\tau=\tau^{\bigstar}(\lambda)} = 0.$$

Defining the the tax elasticity of industry emissions $\eta \equiv [dE(\tau)/E(\tau)]\,/[d\tau/\tau] < 0$ and some rearranging gives the expression for $\tau^{\bigstar}(\lambda)$ in the proposition. The further claims on the properties of $\tau^{\bigstar}(\lambda)$ follow by inspection. $\blacksquare$

Proposition 2 shows how the distortion of the political equilibrium tax $\tau^{\bigstar}(\lambda)$ away from the Pigouvian rule $\tau = D'$ is driven by the weighted average pass-through rate across regulated firms, $\sum_{i=1}^{n} \frac{e_i(\tau)}{E(\tau)}[1 - \rho_i(\tau)]$, with weights given by each firm's share in the tax base. Observe that expression for $\tau^{\bigstar}(\lambda)$ does not hinge on the precise functional form of consumers' gross utility $V(\cdot)$.

To understand the result, note that in the Buchanan second-best problem, industry profits effectively measures the extent of the market-power distortion while in the Grossman-Helpman lobbying problem, industry profits drive the incentive to make political contributions. Proposition 1 told us that, in the GLM class of models, firm-level pass-through pins down firm-level profit impacts—and so the industry-level analog in Proposition 2 is driven by a weighted average of pass-through rates across firms.

Intuitively, lower pass-through means that a firm contracts output more strongly, creating greater deadweight losses and suffering larger profit losses, thus pushing $\tau^{\bigstar}(\lambda)$ downwards—more strongly for large high-emissions firms. Relatedly, where the government is more open to lobbying, and the industry is opposed to the regulation, then this pushes the political equilibrium tax downwards.[61]

Proposition 2 shows how the impact of the Grossman-Helpman lobbying effect ($\lambda > 0$) is, in fact, driven by exactly the same forces as those underlying the Buchanan market-power effect. It also generalizes existing analysis of second-best emissions taxes ($\lambda = 0$) to all models consistent with the GLM, and clarifies the underlying economic intuition in terms of pass-through.

[61] Under perfect competition, each firm's pass-through is 100% so the Pigouvian rule $\tau = D'$ applies—and there is no political lobbying (even if $\lambda > 0$) since no firm is making any profit.

5.2 An empirical estimate for US airlines

We now estimate the political equilibrium carbon price for US airlines. More precisely, we consider a domestic US policymaker chooses her utility-maximizing level of complete regulation for all US airlines, cognizant of the presence of market power in the airline market and under influence of political lobbying by airlines.

We make the result from Proposition 2 operational in four steps. First, we assume a constant social cost of carbon of $D'(\cdot) = \$50/\text{tCO}_2$ (Nordhaus, 2017). Second, we use our previous estimates to obtain the industry-wide average pass-through rate across regulated firms, $\sum_{i=1}^{n} \left(\frac{e_i}{E}\right) \rho_i = 0.78$. As discussed above, this implies that industry profits fall significantly with a tighter carbon price. Third, for the lobbying parameter λ, we turn to the literature. Goldberg & Maggi (1999) were the first to empirically estimate this parameter, finding $\lambda = 0.02$ for the US. McCalman (2004) and Mitra, Thomakos & Ulubasoglu (2002) obtain similar results for Australia and Turkey respectively while Gawande & Bandyopadhyay (2000) find a much higher estimate of $\lambda = 0.5$. Based on these findings, we take $\lambda = 0.1$ as our baseline.

Fourth, analogously to our pass-through analysis, we estimate the *carbon*-tax elasticity of emissions η by estimating the elasticity of industry-level carbon emissions E_t with respect to historical *fuel* prices using the following specification:

$$\log(E_t) = \sum_{m=0}^{3} \eta_m \log(f_{t-m}) + X_t'\beta + \epsilon_t \tag{10}$$

where f_t is the quarterly average of the spot fuel price, and X_t are (emissions-weighted) averages of the same covariates used in our pass-through regressions. Since each gallon of jet fuel produces a constant 0.00957 tons of CO_2, the elasticity of carbon emissions with respect to fuel price is identical to its elasticity with respect to the carbon price.

For consistency, we again estimate this regression using our sample of continuously operated routes over 2002 to 2014 and allow for adjustment to take up to a year (i.e., 4 quarters). We find an "equilibrium" elasticity of $\eta \equiv \sum_{m=0}^{3} \eta_m = -0.16$ (with confidence interval -0.16 ± 0.08). This confirms the intuition that a carbon price is successful at decreasing industry-wide emissions—but also that the elasticity is only modest. Our estimate is at the lower end of the range of elasticities reported in Fukui & Miyoshi (2017).

Using these various baseline parameter values in the formula of Proposition 2 yields a baseline estimate of $\tau^{\star} = \$18.87/\text{CO}_2$ for the political equilibrium carbon tax. This lies more than 60% below the standard Pigouvian tax set at the social cost of carbon ($D'(\cdot) = \$50/\text{tCO}_2$). The gains from the lower equilibrium tax accrue to the large legacy carriers who have limited carbon cost pass-through. By contrast, our analysis suggests that Southwest would actually prefer the carbon tax to be set at the Pigouvian level.

		Lobbying influence (λ)			
		0	0.1	0.2	0.5
	–0.06	\$13.33 (100%)	\$11.63 (96%)	\$10.31 (92%)	\$7.69 (87%)
Price elasticity of emissions (η)	–0.16	\$21.05 (100%)	\$18.87 (93%)	\$17.09 (88%)	\$13.33 (79%)
	–0.26	\$26.09 (100%)	\$23.81 (91%)	\$21.90 (85%)	\$17.65 (74%)

Table 6: Estimates of political equilibrium carbon price for US airlines.

Notes: Uses Proposition 2 based on unified model of Buchanan (1969) market-power effect and Grossman-Helpman (1994) lobbying effect. Social cost of carbon (SCC) set at \$50/tCO_2. Share of Buchanan effect in difference between SCC and equilibrium carbon price in parenthesis.

Table 6 shows sensitivity analysis for the value of $\tau^\star$ when varying the emissions elasticity η and the lobbying parameter λ. Intuitively, when emissions are more responsive to the carbon price, this amplifies its environmental benefits—pushing $\tau^\star$ up. As expected, greater government openness to lobbying pushes $\tau^\star$ down. For a wide range of parameter values, $\tau^\star$ is below half of the SCC.[62]

We can decompose this "shortfall" into two underlying distortions. The Buchanan market-power effect is logically prior in that it can exist without any political lobbying (i.e., where $\lambda = 0$) but the reverse is not true.[63] Thus switching off the lobbying channel, we find $\tau^\star(\lambda = 0) = \$21.05/\mathrm{tCO}_2$. In our baseline estimate, we therefore attribute 93% of the "shortfall" to market power and the remaining 7% to lobbying, given the presence of market power. Table 6 provides sensitivity analysis showing that the share of the Buchanan effect typically exceeds 80%.

In sum, this analysis suggests that the political economy of carbon pricing for an industry with strong product differentiation may be very different from well-researched sectors such as electricity and cement. It may also help explain why aviation has, so far, been a climate laggard. Looking ahead, it suggests that policies to address market power may be able to complement policies to address environmental externalities.

6 Conclusion

We have developed the GLM—a new, simple, flexible reduced-form model of imperfect competition that nests many existing oligopoly models as special cases. We showed that,

[62] It is worth stressing that this finding is not driven by "incomplete" regulation where a carbon price applies only to a subset of firms competing in an industry (Fowlie, Regaunt & Ryan 2016).

[63] In our setting, without any market power, there are no profits and hence nothing to lobby over and so the Grossman-Helpman effect is zero.

within the GLM, firm-level cost pass-through *alone* is a sufficient statistic for the profit impact of regulation on individual firms. Compared with the existing pass-through literature, the GLM relies heavily on supply linearity but allows near-arbitrary firm heterogeneity and does not require additional information on conduct parameters or mark-ups.

We have presented *ex ante* empirical estimates of the impacts of future carbon pricing for US airlines. We found considerable *intra-industry* heterogeneity in pass-through between legacy and low-cost carriers, driven by differences in product portfolios, cost structures and consumer demand. Pass-through is heterogeneous even for a cost shock that hits all firms equally. From a policy perspective, we therefore expect these carrier types to have very different incentives to embrace climate regulation. Our estimated political equilibrium carbon price suggests that the combination of market power and lobbying may be able to explain why aviation has been a climate laggard.

We hope that the GLM will prove useful in other contexts in industrial organization, public economics, international trade, and networks. In this paper, we have shown its value in radically simplifying incidence analysis. More broadly, the GLM lends itself to large-scale estimation both in a *single-industry* context characterized by complex firm heterogeneity demand, costs and conduct and for *cross-industry* analysis that seeks to apply a consistent structure across many different markets. Relative to widely-used structural empirical modeling, its comparative advantage lies in lower complexity and greater transparency in addressing a narrower range of questions.

Appendix A: Extensions to the GLM

We here present three theoretical extensions of the GLM. The first is a simple model of oligopolistic second-degree price discrimination in which customers who buy earlier pay less. The second is a special case of the upgrades approach to multiproduct quality competition (Johnson & Myatt 2003, 2006). The third presents an alternative form of emissions abatement via an "end of pipe" technology that cleans up production *ex post*. In all three settings, our main result from Proposition 1(a), on firm-level cost pass-through as a sufficient statistic for the profit impact of regulation, is preserved.

A.1 Oligopolistic second-degree price discrimination

In practice, consumers often pay different prices for the same good; for example, airlines price discriminate by selling same-class tickets at different prices depending on how far in advance a customer buys. Building on Hazledine (2006), we here show how Proposition 1(a) extends to a linear-symmetric Cournot oligopoly with price dispersion. Consumers have unit demand for a homogeneous product (e.g., an economy flight from A to B), with a distribution of $V = 1-X$ (so X^{th} keenest consumer has value $1-X$). There are H price buckets where class h is priced at $p_h = 1 - X_1 - X_2 - ... - X_h$, where X_h is the number of units sold in class h. Each firm i chooses how much of each price bucket to supply $\{x_{ih}\}_{h=1}^{H}$. Suppose that A1–A3 from the GLM are met, and for simplicity assume that marginal costs $k(\tau) = c(\tau) + \tau z(\tau)$ are symmetric across firms (so all firms are subject to regulation). Firm i's profits $\Pi_i = \sum_{h=1}^{H} \Pi_{ih} = \sum_{h=1}^{H} [p_h - k(\tau)]x_{ih}$ so again defining the profit-neutrality factor via $\Delta\Pi_i(\tau) = -\gamma_i \tau e_i(0)$, it follows that:

$$\gamma_i(\tau) = \frac{-\Delta\Pi_i(\tau)}{\tau z(0)x_i(0)} \Longrightarrow \gamma_i(0) = \frac{1}{z(0)x_i(0)} \left(-\sum_{h=1}^{H} \frac{d\Pi_{ih}(\tau)}{d\tau}\bigg|_{\tau=0} \right),$$

where $x_i \equiv \sum_{h=1}^{H} x_{ih}$ and we again focus on the "small τ" case $\gamma_i(0)$. Using results from Hazledine (2006, *eq* 11 & 13), it is easy to check that A4 is met for each price bucket in the symmetric Nash equilibrium, that is, $[p_k - k(\tau)] = \varphi_h x_{ih}$ where $\varphi_h = (N^H + ... + N + 1)/N^h$. So we have that

$$\begin{aligned}
\sum_{h=1}^{H} \frac{d\Pi_{ih}(\tau)}{d\tau}\bigg|_{\tau=0} &= 2\sum_{h=1}^{H} \frac{1}{\varphi_h}[p_h - k(0)] \left[\frac{dp_h}{d\tau} - \frac{dk}{d\tau}\right]_{\tau=0} \\
&= 2z(0)x_i(0)\sum_{h=1}^{H} \frac{x_{ih}(0)}{x_i(0)} (\rho_h - 1),
\end{aligned}$$

where $dk/d\tau|_{\tau=0} = z(0)$, again by A2 and A3, and $\rho_h \equiv (dp_h/d\tau)/(dk/d\tau)$ is the pass-through rate for price bucket h. This means that the profit-neutrality factor can be written as:

$$\gamma_i(0; H) = 2\sum_{h=1}^{H} \frac{x_{ih}(0)}{x_i(0)} (1 - \rho_h) = 2[1 - \rho_{\text{ave}}(H)],$$

where $\rho_{\text{ave}}(H) \equiv \sum_{h=1}^{H} \frac{x_{ih}(0)}{x_i(0)} \rho_h$ is the *average* pass-through rate across the H price buckets. Hazledine (2006, Proposition 2) shows that price discrimination does not affect the average price paid by consumers, that is, $\sum_{h=1}^{H} \frac{X_h}{X} p_h = p_{\text{ave}}(H) = p_{\text{ave}}(1)$ for all H, where $p_{\text{ave}}(1)$ is the standard case of a uniform (single-bucket) price. This implies that $\frac{d}{d\tau} p_{\text{ave}}(H) = \frac{d}{d\tau} \sum_{h=1}^{H} \frac{X_h}{X} p_h = \frac{d}{d\tau} p_{\text{ave}}(1)$ for all H. Finally, the linear model structure implies that $\frac{d}{d\tau}\left(\frac{X_h(\tau)}{X(\tau)}\right) = 0$, that is, higher cost decreases in equal proportion the output of each price bucket, so the ratio between output of bucket h and total output across all H buckets remains unchanged. Combining these findings shows that:

$$
\begin{aligned}
\frac{d}{d\tau} p_{\text{ave}}(H) &= \frac{d}{d\tau} \sum_{h=1}^{H} \frac{X_h}{X} p_h = \sum_{h=1}^{H} \frac{X_h}{X} \frac{dp_h}{d\tau} = z(0) \sum_{h=1}^{H} \frac{X_h}{X} \rho_h \\
&= z(0) \sum_{h=1}^{H} \frac{x_{ih}}{x_i} \rho_h = z(0)\rho_{\text{ave}}(H) = z(0)\rho_{\text{ave}}(1) = \frac{d}{d\tau} p_{\text{ave}}(1) \text{ for all } H,
\end{aligned}
$$

which uses firm symmetry $x_{ih}/x_i = X_h/X$. We conclude that:

$$
\gamma_i(0; H) = 2[1 - \rho_{\text{ave}}(H)] = 2[1 - \rho_{\text{ave}}(1)] = \gamma_i(0; 1) \equiv \gamma_i(0) \text{ for all } H,
$$

so that irrespective of the number of price buckets, it is the *average* pass-through that drives the profit impact. As in the uniform-price model, a substantive restriction is that the number of firms n does not vary with regulation τ; in addition, this here also holds for the number of price buckets H. Price discrimination raises (lowers) the prices paid by high-value (low-value) consumers but leaves the average price unchanged.[64] This benefits firms because sales expand to otherwise excluded low-value consumers while the average profit margin is unchanged. Conversely, cost pass-through to high-value (low-value) consumers declines (rises) but *on average* is unchanged as the number of price buckets changes.

A.2 Multiproduct competition and the upgrades approach

We here use a simplified version of the upgrades approach to multiproduct quality competition (Johnson & Myatt 2003, 2006; see also the useful exposition in Johnson & Rhodes 2018) to extend our main GLM result to a setting in which firms offer multiple products . We assume that the industry consists of $n \geq 2$ firms each offering two product qualities: low-quality q_1 and high-quality q_2, where $\Delta q \equiv (q_2 - q_1) > 0$. For a regulation τ, we assume that A1–A3 from the GLM are satisfied for each firm and each product quality. Slightly adjusting notation, firm i's unit cost of producing the low-quality product is $k_1^i(\tau) = c_1^i(\tau) + \tau z_1^i(\tau)$ while it is $k_2^i(\tau) = c_2^i(\tau) + \tau z_2^i(\tau)$ for the high-quality product, where $\Delta k^i(\tau) \equiv [k_2^i(\tau) - k_1^i(\tau)] > 0$. We consider a special case that yields linear demand structures for both products. A consumer of type θ has a multiplicative willingness-pay

[64] A rough intuition is that firms view the problem in two steps: first, choosing an aggregate level of output—and hence average price, and, second, choosing how to vary output by bucket—and hence the set of discriminatory prices.

of $v(\theta, q) = \theta q$ for a single unit with quality $q \in \{q_1, q_2\}$. There is a unit mass of potential buyers with uniformly distributed types $\theta \sim U[0,1]$; let $\theta(z)$ denote the buyer type for which there are z buyers with a higher type θ. Let Z_1^i denote the combined number of low- and high-quality units produced by firm i, let Z_2^i be the number of high-quality units, so that $(Z_1^i - Z_2^i)$ is the number of low-quality units. Let $Z_1 = \sum_{i=1}^n Z_1^i$ be the industry supply of both qualities and $Z_2 = \sum_{i=1}^n Z_2^i$ that of high-quality. The former can be interpreted as the number of "baseline" units while the latter is the number of quality "upgrades". The marginal buyer of the low-quality product is indifferent to instead buying nothing and has type $\theta(Z_1)$ so the price of the low-quality product satisfies $p_1(Z_1) = v(\theta(Z_1), q_1) = (1 - Z_1)q_1$. The marginal buyer of an upgrade from low to high quality is indifferent between buying either and has type $\theta(Z_2)$ so the price of upgrading to the high-quality product satisfies $p_2(Z_2) = v(\theta(Z_2), q_2) - v(\theta(Z_2), q_1) = (1 - Z_2)\Delta q$. The total price of the high-quality product is the baseline low-quality price plus the upgrade price, $p_1 + p_2$. Observe that both demand curves $p_1(Z_1)$ and $p_2(Z_2)$ are linear. Firms compete on quantities à la Cournot. Firm i produces Z_1^i baseline units and $Z_2^i \leq Z_1^i$ upgrades (alternatively, $(Z_1^i - Z_2^i)$ low-quality and Z_2^i high-quality units). Firm i's profits are therefore given by:

$$\begin{aligned}\Pi^i(\tau) &= \left[p_1(Z_1) - k_1^i(\tau)\right](Z_1^i - Z_2^i) + \left[p_1(Z_1) + p_2(Z_2) - k_2^i(\tau)\right] Z_2^i \\ &= \left[p_1(Z_1) - k_1^i(\tau)\right] Z_1^i + \left[p_2(Z_2) - \Delta k^i(\tau)\right] Z_2^i,\end{aligned}$$

where the second term reflects the number of upgrades Z_2^i sold incurring an upgrade cost Δk^i. It follows that each firm can separately choose Z_1^i and Z_2^i (subject to $Z_1^i \geq Z_2^i$). As firm i sells both product qualities, using the linear demand curves $p_1(Z_1)$ and $p_2(Z_2)$ from above, its two first-order conditions for Z_1^i and Z_2^i (taking as given rivals' outputs) are given by:

$$\begin{aligned}0 &= \left[p_1(Z_1) - k_1^i(\tau)\right] - Z_1^i q_1 \\ 0 &= \left[p_2(Z_2) - \Delta k^i(\tau)\right] - Z_2^i \Delta q.\end{aligned}$$

This shows that the GLM's A4 assumption on linear supply behaviour holds for each firm and each product; in particular, $\psi_1^i = q_1^{-1}$ for the low-product and $\psi_2^i = (\Delta q)^{-1}$ for the upgrade (both $\forall i$). Hence Proposition 1 from the main text applies and cost pass-through is a sufficient statistic for the profit impact of regulation on each firm-product. Similar to the preceding extension with second-degree price discrimination, a substantive restriction is that the quality levels q_1, q_2 offered by firms do not vary with regulation τ. In sum, the upgrades approach to multiproduct competition has the key feature that the baseline units and upgrades are neither substitutes nor complements; given A1–A3 and linear demand structures, the GLM's A4 holds for each individual product category. Finally, for our empirical application to US airlines, the upgrades model provides a microfoundation for the widespread approach in the literature of analyzing economy-class ticket pricing (as the baseline product) in a separable way from business- or first-class tickets (viewed as product upgrades).

A.3 Emissions abatement with an end-of-pipe technology

We here show that the result of Proposition 1(a) also holds if firm i instead has access to an "end-of-pipe" abatement technology. In the case of climate policy, this could be carbon capture and storage (CCS)—which has the potential to be applied in power generation as well as a range of industrial sectors. Other examples of end-of-pipe technologies are scrubbers on smokestacks, catalytic converters for cars, and various technologies for the treatment of industrial waste water. Formally, we replace A2–A3 with the following alternative assumptions:

A2′. (*Cost-minimizing emissions*) Firm i chooses its inputs, including emissions e_i, optimally so as to minimize its total costs $C_i(x_i, e_i) + \tau e_i + \varphi_i(\overline{z}_i x_i - e_j)$ of producing output x_i, where $\overline{z}_i \equiv z_i(0)$ is its fixed emissions intensity (emissions per unit of output) and the investment cost of the end-of-pipe technology satisfies $\varphi_i'(\cdot), \varphi_i''(\cdot) > 0$ with $\varphi_i(0) = \varphi_i'(0) = 0$.

A3′. (*Constant returns to scale*) Firm i's optimized total costs at an emissions price τ are affine in output $C_i(x_i, e_i) + \tau e_i = k_i(0)x_i + \varphi(\overline{z}_i x_i - e_i)$, with unit cost $k_i(0) = c_i(0) + \tau\overline{z}_i$. The emissions intensity of production remains fixed at $\overline{z}_i$ but, following production, the

firm can clean up some or all of the resulting emissions at a fixed cost. Firm i's optimal choice of emissions satisfies $\partial\Pi_i/\partial e_i = -\tau + \varphi_i' = 0$ and so $e_i = \overline{z}_i x_i - \varphi_i'^{-1}(\tau)$ for given x_i. Hence optimal profits can be written as:

$$\Pi_i(\tau) = [p_i - c_i(0) - \tau\overline{z}_i]\, x_i + \underbrace{\left[\tau\varphi_i'^{-1}(\tau) - \varphi_i(\varphi_i'^{-1}(\tau))\right]}_{\equiv \upsilon_i(\tau)}.$$

The optimal value $\upsilon_i(\tau) \geq 0$ of the end-of-pipe abatement technology has the following properties:

$$\begin{aligned}
\upsilon_i'(\tau) &= \varphi_i'^{-1}(\tau) + \left[\tau - \varphi_i'(\varphi_i'^{-1}(\tau))\right] \frac{d}{d\tau}\varphi_i'^{-1}(\tau) = \varphi_i'^{-1}(\tau) \geq 0 \\
\upsilon_i''(\tau) &= \frac{d}{d\tau}\varphi_i'^{-1}(\tau) = \frac{1}{\varphi_i''\left(\varphi_i'^{-1}(\tau)\right)} > 0.
\end{aligned}$$

Similar to the value of exercising an option, $\upsilon_i(\tau)$ is increasing and convex in the underlying source of variation—here the carbon price. End-of-pipe abatement has no effect on i's product-market strategy; this is identical to the GLM above, given unit cost

$k_i(0) = c_i(0) + \tau \bar{z}_i$. For a small carbon price, $v_i'(0) = 0$ because $\varphi_i'(0) = 0$; the optimal value of the abatement technology is of second order—and hence Proposition 1(a) applies for small τ. Of course, changes in the abatement technologies available to j's *rivals* can alter i's own profit impact; again, the point is that, given A1, A2$'$, A3$'$, A4, any such differences will be appropriately reflected in i's *own* pass-through rate.

Appendix B: Construction of airline data

Ticket price p_{ijt} we obtain from the cleaned DB1A data provided by Severin Borenstein (the raw DB1A data, along with all the data below, are from the Bureau of Transportation Statistics). We drop any non-direct tickets for ijt, and then convert the nominal prices to real 2014Q4 USD using St. Louis Fed CPI data (as we do with all monetary variables).

Per-passenger fuel cost k_{ijt} is constructed as follows, with the raw variable names given parentheses. First we use the Form 41 (Schedule P-5.2) dataset, which contains carrier-aircraft-time specific fuel costs (fuel_fly_ops), which we denote k_{ilt}. Following O'Kelly (2012), we assume the fuel used to fly route j is a linear function of distance d_j with a non-zero intercept: $k_{ilj} = b^0_{ilt} + b^1_{ilt} d_{ilj}$. The fixed cost comes from the fuel used in take-off and landing, and any airport related activities; the variable cost is the 'miles per gallon' fuel consumption at cruising altitude. The fuel use data we have do not allow us to identify both the slope and the intercept, so we use an average value for their ratio taken from EEA (2016): we set the ratio $\frac{b^0_{ilt}}{b^1_{ilt}} = 131$ for all ilt, meaning take off and landing uses the same fuel as cruising 131 miles. Next we use the T-100 Domestic Segment to assign aircraft to routes. We construct the share $\alpha_{ijt}(l)$ of carrier i's passengers on route j at time t that travelled on aircraft type l (aircraft_type). We use total 'effective distance' flown by each aircraft type l on each route j, $\tilde{d}_{iljt} = (\frac{b^0}{b^1} + \text{distance}_j) \times \text{dep\_performed}_{ijlt}$, so that $\alpha_{ijt}(l) = \frac{\tilde{d}_{ijlt}}{\sum_l \tilde{d}_{ijlt}}$. Using these shares we construct the weighted average fuel cost $k_{ijt} = \sum_l \alpha_{ijt}(l) k_{ilt}$.

Labour and maintenance cost index c_{jt} is constructed from Form 41 (Schedule P-5.2). We take, for each ilt, total flying operating costs (tot_fly_ops) plus total maintenance costs (tot_dir_maint) minus fuel costs (fuel_fly_ops). We then construct a weighted average value for each ijt using the weights $\alpha_{ijt}(l)$ described above. Finally, we transform the carrier-route-time specific costs (which could be subject to endogeneity via their denominator), into a carrier-time index of costs. This is done by dividing total costs by total passengers, for each carrier-time. We normalise to the 2000Q1 value for American Airlines.

GDP growth g_{jt} is constructed with data from the Federal Reserve Bank of St. Louis. Using state-level GDP data, for each route j we take the average of the states in which each of the origin and destination airports are located. For the period 2002-2004 we interpolated the annual data as quarterly data is not availble.

Number of competitor firms n_{jt} we construct from the DB1A data. We define competitors to include all routes that serve the same city-city market as route j. For example LAX-SLC is a competitor product to SNA-SLC because LAX and SNA both serve the city of Los Angeles. Using the Bureau of Transportation Statistics' definition of a market (origin_city_market_id and dest_city_market_id), we count all carriers serving that market with at least 1,000 passengers in a quarter.

Number of potential entrants n_{jt}^p is constructed from the DB1A. Using the above city markets, we count the number of carriers operating from each end of route j at time t, and subtract the number of carriers operating a direct service on that market route. This gives n_{jt}^p.

Appendix C: Empirical results

C.1 Regression result for a single route

The estimation results for equation (8), the first stage of the 2SLS estimation, give F-statistics for each of the four lags of the endogenous fuel cost of 57, 40, 86, 213 respectively. The four Sanderson-Windmeijer chi-squared Wald statistics lie in the range 83-225, suggesting no underidentification. The four Sanderson-Windmeijer F-statistics are in the range 11.4-31.2, suggesting the IV bias relative to OLS is likely to be small. These results are not surprising, given the strong correlation between airlines' fuel costs and spot fuel prices summarised in Figure 2a. The second stage regression, equation (7), is reported in Table 7. Significance levels for the covariates are generally quite low, but this is to be expected given the small number of observations and relatively large number of regressors. The coefficients all have the signs we would expect, with the exception of n, which has very little variation for this route (see Figure 2).

C.2 Full baseline results

Table 8 gives the weighted averages over all carrier-routes ij of the 2SLS estimates of the coefficients in baseline specification (7). The signs of the coefficients are all as expected, and the results are generally highly significant. See Section 4.6 for further discussion.

C.3 Estimates for individual legacy carriers

Table 9 shows the descriptive statistics by individual carrier. There is heterogeneity within the legacy carriers, with Alaska and Hawaiian in particular flying shorter and therefore

Pass-through	1.38
	(0.32)
No. firms	2.05
	(3.26)
No. potential entrants	-2.11
	(2.03)
Labour & maintenance cost index	166.81
	(99.12)
GDP growth	537.72
	(281.76)
Quarter 1	-3.87
	(7.87)
Quarter 2	5.55
	(4.54)
Quarter 3	15.81
	(5.58)
Constant	113.99
	(17.20)
R squared	0.68
No. observations	52

Table 7: Estimation results on pass-through for Southwest's route from Phoenix (PHX) to San Antonio (SAT).

Notes: Results for 2SLS estimation of (7) for PHX-SAT. The dependent variable is ticket price. The pass-through is the sum of coefficients of 4 quarterly lags. Corresponds to the results reported in Table 3 in the main text. Standard errors in parentheses.

	Southwest	Legacy	All
Pass-through	1.48	0.55	0.78
	(0.03)	(0.06)	(0.05)
No. firms	-1.91	-7.08	-5.77
	(0.37)	(0.84)	(0.65)
No. potential entrants	-1.13	-1.13	-1.13
	(0.15)	(0.42)	(0.32)
Labour and maintenance cost index	122.66	97.88	104.17
	(8.69)	(6.53)	(5.51)
GDP growth	173.85	93.21	113.68
	(18.44)	(53.27)	(40.91)
Quarter 1	-5.75	-7.97	-7.41
	(0.53)	(1.69)	(1.30)
Quarter 2	4.32	10.94	9.26
	(0.48)	(1.23)	(0.95)
Quarter 3	-1.71	12.77	9.10
	(0.50)	(1.47)	(1.13)
R squared	0.67	0.45	0.56
No. routes	212	403	615
No. observations	11,024	20,956	31,980

Table 8: Pass-through estimates for Southwest and the legacy carriers, with covariates.

Notes: Results for 2SLS estimation of (7) using all carrier-routes. The dependent variable is ticket price. The pass-through is the sum of coefficients of 4 quarterly lags. Corresponds to the results reported in Table 3 in the main text, but with coefficient covariates included.

cheaper routes. This is can be explained by their emphasis on services within the US's two particularly remote States, combined with services from the mainland US to these states. These airlines are a small fraction of the total legacy portfolio, however, and the four largest legacy carriers are markedly more homogenous. Despite these two somewhat anomalous small legacy carriers, Southwest is still the outlier. Table 10 shows the carrier average pass-through rates, broken down by carrier within the legacy group. Southwest's pass-through is significantly above 1 by a considerable margin, whereas all the legacy carriers have a pass-through point estimate below 1. There is, however, considerable variation within the group. At the lower end, United's pass-through has a negative point estimate, but it is not statistically significantly different from 0. At the upper end, American, Hawaiian and US Airways all have pass-through rates that are not significantly different from 1 at 5%. Nonetheless, Southwest clearly remains significantly higher than the legacy carriers, whose pass-through rates can plausibly be claimed to likely lie in the

	WN	AA	AS	DL	HA	UA	US
Price ($)	157.31	226.29	205.46	230.86	166.68	245.56	240.44
Fuel cost ($)	29.22	54.52	43.36	47.20	41.54	55.32	42.15
Distance (miles)	688	1,163	726	1,041	1,110	1,277	957
Emissions (tCO_2)	0.13	0.24	0.18	0.19	0.17	0.22	0.18
Emissions cost ($)	6.70	12.04	9.13	9.39	8.33	11.15	9.06
Passengers (000s)	195	159	158	155	331	141	127
No. firms	3.28	3.79	2.57	3.35	2.78	4.65	3.05
Fraction seats filled	0.72	0.79	0.70	0.81	0.81	0.81	0.79
Revenue ($ million)	24.76	31.46	24.82	29.36	35.12	29.46	24.19
Revenue in sample	0.42	0.39	0.41	0.26	0.40	0.45	0.27
No. routes	212	111	35	90	10	101	56
No. observations	11,024	5,772	1,820	4,680	520	5,252	2,912

Table 9: Descriptive statistics for individual legacy carriers.

Notes: Passenger numbers and revenue are year averages over the 2002Q1-2014Q4 sample period. Emissions cost calculated using $50/$tCO_2$ carbon price. WN = Southwest, AA = American Airlines, AS = Alaska Airlines, DL = Delta, HA = Hawaiian Airlines, UA = United Airlines, US = US Airways.

interval $(0, 1)$.

Appendix D: Robustness of empirical results

In this section, we report regression results for modified or extended specifications, or for the baseline specification estimated using subsets of the full sample. Our purpose is to show some further results indicative of factors important to pass through, and to show that the numerical results reported in Section 4 are robust to the relevant alternative setups.

D.1 OLS estimates

For completeness and to provide a benchmark, we report the results of estimating equation (7) using OLS, rather than the baseline 2SLS approach described in Section 4. The result, reported in Table 11, is pass-through rates of 1.34 and 0.43 for Southwest and the legacy carriers respectively. This is suggestive of some common downward bias in OLS, but also confirms that our instruments are not driving the difference between the two airline types.

	WN	AA	AS	DL	HA	UA	US
Pass through	1.48	0.90	0.21	0.79	0.92	-0.09	0.69
	(0.04)	(0.08)	(0.09)	(0.14)	(0.18)	(0.09)	(0.40)
Profit impact (%)	2.95	-0.80	-6.41	-1.39	-0.54	-9.58	-2.31
	(0.22)	(0.69)	(0.70)	(0.94)	(1.31)	(0.76)	(2.93)
No. routes	212	111	35	90	10	101	56
No. observations	11,024	5,772	1,820	4,680	520	5,252	2,912

Table 10: Pass-through estimates and predicted profit impacts for individual legacy carriers.

Notes: Profit impacts based on \$50/tCO_2 carbon price implemented at 2014Q4 using 2014 carrier-route year average emissions intensity. WN = Southwest, AA = American Airlines, AS = Alaska Airlines, DL = Delta, HA = Hawaiian Airlines, UA = United Airlines, US = US Airways.

D.2 Time

We chose the starting period for our baseline results in order to avoid the immediate effects of the terrorist attacks of 9/11 in 2001. However, we have re-estimated specification for a later, and smaller, 10 year time window: 2005Q1-2014Q4. As shown in Table 11, the number of continuously operated routes is slightly larger, but the pass-through rates are statistically unchanged from the full period results, suggesting our choice of time window (which includes both good times and bad for US airlines) is not driving the results. We tried all possible time periods that contained at least 10 consecutive years within the overall window 2002-2014, and in all cases the qualitative results in the baseline continued to hold.

D.3 Entry and exit

When estimating pass-through rates using variation over time, we implicitly require market structure to be stable over time. In reality, however, there is considerable exit and entry on many routes over the period of our sample, which could correlate with fuel cost. This was our motivation for including both number of rivals and number of potential entrants in all our regressions, such that any effects on the level of the price are controlled for. However, we here explore whether entry and exit may be an issue in two further ways. First, we allow the pass-through rate to be a linear function of the number of competitors. We take a simple approach, which we repeat several variants of below, by including an interaction between a variable of interest v_{ijt} and the fuel cost, k_{ijt}, so that

	Southwest	Legacy
(a) Baseline (2SLS)	1.48	0.55
	(0.03)	(0.06)
	212	*403*
(b) OLS	1.34	0.43
	(0.03)	(0.04)
	212	*403*
(c) Late period: 2005-2014 only	1.50	0.62
	(0.06)	(0.06)
	229	*413*
(d) n-interaction	1.45	0.64
	(0.04)	(0.07)
	212	*403*
(e) Baseline with $\Delta n = 0$	1.54	0.66
	(0.12)	(0.19)
	24	*17*
(f) Baseline with $\Delta n \leq 1$	1.63	0.82
	(0.08)	(0.12)
	50	*57*
(g) Fixed effects specification	1.31	0.57
	(0.05)	(0.06)
	212	*403*
(h) Log specification	0.21	0.15
	(0.01)	(0.01)
	212	*403*

Table 11: Robustness analysis of pass-through estimates across carrier types.

Notes: Results using alternative 2SLS estimates of specification (7), different subsets of the full sample, and other alternative specifications. Regression (a) is the same as Table 3 in the main text. Regressions (b)-(h) are described in Appendix D. Standard errors in parentheses, number of routes in italics.

the specification becomes:

$$p_{ijt} = \sum_{m=0}^{3} \rho_{ij}^{m} k_{ij,t-m} + \sum_{m=0}^{3} \xi_{ij}^{m} k_{ij,t-m} \cdot v_{ij,t-m} + X'_{ijt}\beta_{ij} + \epsilon_{ijt} \tag{11}$$

The equilibrium rate of pass-through is now $\rho_{ij} = \sum_{m=o}^{3} \rho_{ij}^{m} + \sum_{m=o}^{3} \xi_{ij}^{m} \cdot v_{ij}$. We can test whether the variable v impacts the pass-through rate by looking at the statistical significance of $\xi_{ij} = \sum_{m=0}^{3} \xi_{ij}^{m}$. To test whether entry impacts on the pass-through, we set $v_{ijt} = n_{ijt}$, the number of firms operating on a route, and evaluate the pass-through rate in 2014 using $n_{ij,t=2014}$, the average number of competitor firms in 2014. The resulting pass-through rates, shown in Table 11 row (d), are stasticially indistinguishable from the baseline case. Table 12 row (a) shows the value of the interaction coefficient, which is not statistically different from zero. The second way we explore the impact of entry and exit is to recalculate average pass-through rates, but only for subsets of routes that see zero or little entry and exit over the period. The results are shown in Table 11, rows (e) and (f), and demonstrate a reasonably stable pass-through rate for the different types of route. We would not expect these to be identical; the $\Delta n = 0$ routes, for example, are all monopolies. We consider this more of a check that the results don't appear to be being driven by routes with either particularly stable or unstable numbers of competitors. Southwest's pass-through rates are very stable across the different subsets, and statistically indistinguishable. The legacy carriers are also statistically indistinguishable at 10%, though not at 5%. Similar results are obtained when looking at stability in the number of potential entrants, n^p.

D.4 Asymmetric pass-through

Consumers and regulators often suspect that prices rise faster and further after an increase in firm costs than they fall after a fall in costs, a phenomenon known as "rockets and feathers". We test for the presence of asymmetric pass-through by setting the interaction variable v_{ijt} in specification (11) equal to a dummy variable equal to 1 if $p_{ijt} > p_{ij,t-1}$, and equal to 0 otherwise. In order to produce results directly comparable to Peltzman (2000), we run the regression in first differences rather than levels. These results are therefore not directly comparable with our baseline results. We find the equilibrium coefficient on the interaction term is significantly different (at 5%) from 0 on 9.9% of carrier-routes, with no considerable difference between between Southwest (8.0%) and the legacy carriers (10.9%). Hence, cost pass-through seems to generally be symmetric for airlines in this setting. This puts air travel in the minority: Peltzman (2000) found around two thirds of a broad range of goods exhibit asymmetric cost pass-through.

D.5 Fixed effects estimation

As argued in Section 4, running individual regressions rather than Fixed Effects estimation is the correct approach in this setting, as it does not impose homogeneity restrictions on the pass-through rates or the effect of covariates on each airline's different routes. For completeness, we report the results of a 2SLS fixed effects estimation of specification (7) in Table 11. Although the pass-through result for Southwest is a little lower than in the baseline case, our overall qualitative findings are similar in this miss-specification.

D.6 Log specification

Section 4 reports pass-through rates in levels. These are the relevant input for our theoretical results, and this is the standard approach in much of the pass-through literature. However, others areas of the literature, for example, on exchange rate pass-through, estimate pass-through elasticities. We therefore estimate a version of equation (7) using logs of price and fuel cost. The pass-through rates, reported in row (h) of Table 11, lie in the interval $(0, 1)$, as expected for an elasticity measure. Without specifying the level of both price and total marginal cost, we cannot directly compare an elasticity pass-through with the levels pass-throughs in our baseline results. However, the logged results are reassuring in that once again we see a significantly larger pass-through for Southwest than the legacy carriers.

Appendix E: Further empirical results

E.1 Southwest presence

Goolsbee & Syverson (2008), among others, stress the importance of competition with Southwest being a major factor in the pricing decisions of the legacy carriers. We explore this issue by asking whether Southwest's presence, actual or potential, may affect not just the level of the legacy carriers' prices but their pass-through as well. To do this, we first set the interaction variable v_{ijt} in specification (11) equal to a dummy equal to 1 if Southwest competes with route j in time t. Clearly when there is no variation in Southwest's presence over time this effect will not be identified, so the coefficient ξ will be capturing the impact of entry (or exit, but this is less common) by Southwest on the legacy carrier's pass-through. The result in row (d) of Table 12 indicates that Southwest's presence as a competitor significantly reduces the legacy carriers' pass-through rates. In row (e) of Table 12 we show the result of including an additional interaction term in specification (11), a dummy equal to 1 if Southwest is a potential entrant on route j (using the definition from Goolsbee & Syverson 2008, and described in Appendix B). The results suggest that once Southwest has entered a route, its presence has no statistically significant effect on pass-through. However, potential entry by Southwest has a significant and highly negative effect on the pass-through rate. This result compliments and reinforces the findings in Goolsbee & Syverson (2008), suggesting that, separate to a level effect on prices, the threat of entry by Southwest heavily constrains the ability of legacy carriers to pass through cost rises to their customers.

	Southwest	Legacy
(a) No. firms n	0.00	-0.01
	(1.45)	(0.21)
	183	*379*
(b) Volatility	-0.018	-0.010
	(0.001)	(0.001)
	212	*403*
(c) Bankruptcy dummy	–	0.15
	–	(0.03)
	–	*358*
(d) Southwest present dummy	–	-0.24
	–	(0.08)
	–	*209*
(e) Southwest present dummy	–	0.05
	–	(0.20)
	–	*108*
Southwest potential	–	-0.91
	–	(0.36)
	–	*108*

Table 12: Interaction coefficient ξ.

Notes: Results for 2SLS estimation of specification (11), interacting fuel costs with another variable of interest, given in the first column of the table. The resulting ξ shown is the weighted average of the ξ_{ij}, using the same weights as in the main text for ρ. Standard errors in parentheses, number of routes in italics.

E.2 The impact of legacy carrier bankruptcy

Most of the legacy carriers in our sample went through periods of bankruptcy over our period of study: American in 2011-13, Delta in 2005-07, United in 2002-06 and US in 2002-03 and 2004-05. (Southwest did not go bankrupt at any point.) Much has been written on the often weak financial performance of the industry since deregulation (see Borenstein 2011). There is also prior evidence that airlines' financial constraints can lead to price wars (Busse 2002). We here explore the impact of bankruptcy on the cost pass-through of legacy carriers. To see whether bankruptcy has an effect on pass-through we set the interaction variable v_{ijt} in specification (11) equal to a dummy equal to 1 if airline i is bankrupt in quarter t. As shown in Table 12 row (c), the estimate of the average equilibrium coefficient ξ on this interaction term is equal to 0.15 ($\pm$ 0.06), suggesting pass-through rates are higher for carriers in periods of bankruptcy. The finding of higher pass-through is consistent with the intuition that a firm in severe financial distress is less able or willing to absorb cost shocks. This might also be a relevant factor in explaining how these airlines subsequently again came out of bankruptcy.

E.3 Fuel price volatility

Kellogg (2014) finds that expected future oil price volatility has a significant impact on oil drilling investments in Texas, with firms investing less when volatility is high. Though the mechanism would differ, it is possible that airlines' pass-through may also be affected by input cost volatility. This could have important consequences for profits since an emissions trading scheme would lead to carbon price volatility whereas a carbon tax would not. To test the impact of fuel price volatility, we set the interaction term v_{ijt} in specification (11) equal to a simple measure of volatility, the standard deviation, over the quarter, of daily jet fuel prices. The results, in row (b) of Table 12, show that input cost volatility significantly and negatively affects the ability of firms to pass through costs. To give an idea of the economic significance of these effects, we can see by how much the pass through rates would change when fuel cost volatilities are 'low' or 'high', which we take to be in their 10^{th} and 90^{th} percentiles respectively. When volatility is low, pass-through rates change by +0.10 and +0.06 for Southwest and the legacy carriers, respectively; when volatility is high, pass-through rates change by -0.26 and -0.15, respectively. (The asymmetry in this effect comes from the asymmetry of the distribution of quaterly volatility.) This is, to our knowledge, a new result, which suggests that an emissions trading scheme could have a more negative impact on profits than a carbon tax at the same average level.[65]

References

Acemoglu, Daron and Martin Kaae Jensen (2013). Aggregative Comparative Statics. *Games & Economic Behavior* 81, 27–49.

Allaz, Blaise and Vila, Jean-Luc (1993). Cournot Competition, Forward Markets and Efficiency. *Journal of Economic Theory* 59, 1-16.

[65] A caveat is that this finding implicitly assumes that airlines' hedging behavior remains unchanged with higher input cost volatility.

Al-Najjar, Nabil, Sandeep Baliga and David Besanko (2008). Market Forces Meet Behavioral Biases: Cost Misallocation and Irrational Pricing. *RAND Journal of Economics* 39, 214–237.

Anderson, Simon P., André de Palma and Brent Kreider (2001). Tax Incidence in Differentiated Product Oligopoly. *Journal of Public Economics* 81, 173–192.

Ashenfelter, Orley and Robert S. Smith (1979). Compliance with the Minimum Wage Law. *Journal of Political Economy* 87, 333–350.

Ashenfelter, Orley, David Ashmore, Jonathan B. Baker and Signe-Mary McKernan (1998). Identifying the Firm-Specific Cost Pass-Through Rate. Working Paper at the Federal Trade Commission, January 1998.

Atkin, David and Dave Donaldson (2015). Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs. Working Paper at MIT, July 2015.

Azar, José, Martin C. Schmalz and Isabel Tecu (2018). Anti-Competitive Effects of Common Ownership. *Journal of Finance*, forthcoming.

Baumol, William J. and Wallace E. Oates (1988). *The Theory of Environmental Policy.* Cambridge University Press, Cambridge: UK (2nd edition).

Bergquist, Lauren Falcao (2017). Pass-through, Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya. Working Paper at University of Chicago, September 2017.

Bernheim, B. Douglas and Michael D. Whinston (1986). Menu Auctions, Resource Allocation and Economic Influence. *Quarterly Journal of Economics* 101, 1–31.

Berry, Steven (1992) Estimation of a Model of Entry in the Airline Industry. *Econometrica* 60(4): 889–917.

Berry, Steven, James Levinsohn and Ariel Pakes (1995). Automobile Prices in Market Equilibrium. *Econometrica* 63(4) 841–890.

Berry, Steven, and Panle Jia (2010). Tracing the Woes: An Empirical Analysis of the Airline Industry. *American Economic Journal: Microeconomics* 2(3): 1–43.

Bombardini, Matilde (2008). Firm Heterogeneity and Lobby Participation. *Journal of International Economics* 75, 329–348.

Borenstein, Severin (1992). The Evoluation of US Airline Competition. *Journal of Economic Perspectives* 6, 45–73.

Borenstein, Severin (2011). Why Can't US Airlines Make Money? *American Economic Review* 101, 233–237.

Borenstein, Severin and Nancy L. Rose (1994). Competition and Price Dispersion in the US Airline Industry. *Journal of Political Economy* 102, 653–683.

Bovenberg, A. Lans, Lawrence H. Goulder and Derek J. Gurney (2005). Efficiency Costs of Meeting Industry-Distributional Constraints Under Environmental Permits and Taxes. *RAND Journal of Economics* 36, 950–970.

Brander, James A. and Anming Zhang (1990). Market Conduct in the Airline Industry: An Empirical Investigation. *RAND Journal of Economics* 21, 567–583.

Bresnahan, Timothy (1989). Empirical Methods for Industries with Market Power. In: Richard Schmalensee and Robert Willig (Eds.), *Handbook of Industrial Organization*, Volume II, Amsterdam: Elsevier Science Publishers.

Buchanan, James (1969). External Diseconomies, Corrective Taxes, and Market Structure. *American Economic Review* 59, 174–177.

Bulow, Jeremy I. and Paul Pfleiderer (1983). A Note on the Effect of Cost Changes on Prices. *Journal of Political Economy* 91, 182–185.

Bushnell, James B., Howard Chyong and Erin T. Mansur (2013). Profiting from Regulation: Evidence from the European Carbon Market. *American Economic Journal: Economic Policy* 5, 78–106.

Bushnell, James and Jacob Humber (2017). Rethinking Trade Exposure: The Incidence of Environmental Charges in the Nitrogenous Fertilizer Industry. *Journal of the Association of Environmental and Resource Economists* 4, 857–894.

Busse, Meghan (2002). Firm Financial Condition and Airline Price Wars. *RAND Journal of Economics* 33, 298–318.

Carney, Mark (2015). Breaking the Tragedy of the Horizon: Climate Change and Financial Stability. Bank of England, Speech given to Lloyd's of London, 29 September 2015.

Ciliberto, Federico and Elie Tamer (2009). Market Structure and Multiple Equilibria in Airline Markets. *Econometrica* 77, 1791–1828.

Corchón, Luis C. (1994). Comparative Statics for Aggregative Games: The Strong Concavity Case. *Mathematical Social Sciences* 28, 151–165.

Cui, Qiang, Yi-Ming Wei, Chen-lu Yu, Ye Li (2016). Measuring the Energy Efficiency for Airlines Under the Pressure of Being Included into the EU ETS. *Journal of Advanced Transportation* 44, 39–53.

Crooke, Philip, Luke Froeb, Steven Tschantz and Gregory J. Werden (1999). Effects of Assumed Demand Form on Simulated Postmerger Equilibria. *Review of Industrial Organization* 15, 205–217.

Daughety, Andrew (1990). Beneficial Concentration. *American Economic Review* 80, 1231–1237.

De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, Pavcnik (2016). Prices, Markups and Trade Reform. *Econometrica* 84, 445–510.

Draca, Mirko, Stephen Machin and John Van Reenen (2011). Minimum Wages and Firm Profitability. *American Economic Journal: Applied Economics* 3, 129–151.

Einav, Liran and Jonathan Levin (2010). Empirical Industrial Organization: A Progress Report. *Journal of Economic Perspectives* 24, 145–162.

European Environment Agency (2016). EEA Air Pollutant Emission Inventory Guidebook 2016, Chapter 1.A.3, Annex 4. Technical Report, 2016.

Fabra, Natalia and Mar Reguant (2014). Pass-Through of Emissions Costs in Electricity Markets. *American Economic Review* 104, 2972–2899.

Fershtman, Chaim and Kenneth L. Judd (1987). Equilibrium Incentives in Oligopoly. *American Economic Review* 77, 927–940.

Fukui, Hideki, and Chikage Miyoshi (2017). The impact of aviation fuel tax on fuel consumption and carbon emissions: The case of the US airline industry. *Transportation Research Part D* 50, 234–253.

Fowlie, Meredith, Mar Reguant and Stephen P. Ryan (2016). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy* 124, 249–302.

Ganapati, Sharat, Joseph S. Shapiro and Reed Walker (2016). The Incidence of Carbon Taxes in U.S. Manufacturing: Lessons from Energy Cost Pass-Through. NBER Working Paper 22281.

Gaudenzi, Barbara, and Alessandro Bucciol (2016). Jet fuel price variations and market value: a focus on low-cost and regular airline companies. *Journal of Business Economics and Management* 17(6): 977–991.

Gawande, Kishore, and Usree Bandyopadhyay (2000). Is Protection for Sale? Evidence on the Grossman-Helpman Theory of Endogenous Protection. *Review of Economics and Statistics* 82(1), 139–152.

Gerardi, Kristopher S. and Adam Hale Shapiro (2009). Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry. *Journal of Political Economy* 117, 1–37.

Goldberg, Pinelopi K. and Giovanni Maggi (1999). Protection for Sale: An Empirical Investigation. *American Economic Review* 89(5), 1135–1155.

Goolsbee, Austan and Chad Syverson (2008). How Do Incumbents Respond to the Threat of Entry? Evidence from the Major Airlines. *Quarterly Journal of Economics* 123, 1611–1633.

Grossman, Gene M. and Elhanan Helpman (1994). Protection for Sale. *American Economic Review* 84, 833–850.

Häckner, Jonas (2000). A Note on Price and Quantity Competition in Differentiated Oligopolies. *Journal of Economic Theory* 93, 233–239.

Hazledine, Tim (2006). Price Discrimination in Cournot-Nash Oligopoly. *Economics Letters* 93, 413–420.

Hepburn, Cameron J., John K.-H. Quah and Robert A. Ritz (2013). Emissions Trading with Profit-Neutral Permit Allocations. *Journal of Public Economics* 98, 85–99.

ICAO (2016). *2016 Environmental Report*. International Civil Aviation Organisation, Geneva, Switzerland.

IEA (2008) Energy Technology Perspectives 2008: Scenarios and Strategies to 2050. International Energy Agency, Paris, France.

Johnson, Justin P. and David P. Myatt (2003). Multiproduct Quality Competition: Fighting Brands and Product Line Pruning. *American Economic Review* 93, 748–774.

Johnson, Justin P. and David P. Myatt (2006). Multiproduct Cournot Oligopoly. *RAND Journal of Economics* 37, 583–601.

Johnson, Justin P. and Andrew Rhodes (2018). Multiproduct Mergers and Quality Competition. Working paper at Cornell and Toulouse, November 2018.

Kellogg, Ryan (2014). The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling. *American Economic Review* 104(6), 1698–1734.

Kim, Harim (2018). Heterogeneous Impacts of Cost Shocks, Strategic Bidding and Pass-through: Evidence from the New England Electricity Market. Working Paper at University of Mannheim, March 2018.

Kim, E. Han and Vijay Singal (1993). Mergers and Market Power: Evidence from the Airline Industry. *American Economic Review* 83, 549–569.

Klemperer, Paul D. and Margaret A. Meyer (1989). Supply Function Equilibria in Oligopoly under Uncertainty. *Econometrica* 57, 1243–1277.

Kwan, Irene and Daniel Rutherford (2015). Assessment of U.S. Domestic Airline Fuel Efficiency Since 2010. *Transportation Research Record: Journal of the Transportation Research Board* 2501, 1–8.

McCalman, Phillip (2004). Protection for Sale and Trade Liberalization: an Empirical Investigation. *Review of International Economics* 12(1), 81–94.

McCollum, David, Gregory Gould and David Greene (2009). Greenhouse Gas Emissions from Aviation and Marine Transportation: Mitigation Potential and Policies. Pew Center on Global Climate Change.

Menezes, Flavio M. and John Quiggin (2012). More Competitors or More Competition? Market Concentration and the Intensity of Competition. *Economics Letters* 117, 712–714.

Miller, Nathan H., Matthew Osborne and Gloria Sheu (2017). Pass-Through in a Concentrated Industry: Empirical Evidence and Regulatory Implications. *RAND Journal of Economics* 48, 69–93.

Mitra, Devashish, Dimitrios D., Thomakos and Mehmet A. Ulubasoglu (2002). "Protection For Sale" In A Developing Country: Democracy Vs. Dictatorship. *The Review of Economics and Statistics* 84(3), 497–508.

Monderer, Dov and Lloyd Shapley (1996). Potential Games. *Games & Economic Behavior* 14, 124–143.

Morrison, Steven A. (2001). Actual, Adjacent, and Potential Competition: Estimating the Full Effect of Southwest Airlines. *Journal of Transport Economics and Policy* 35(2), 239–256.

Nevo, Aviv (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica* 69, 307–342.

Nordhaus, William D. (2017). Revisiting the Social Cost of Carbon. *Proceedings of the National Academy of Sciences* 114, 1518–1523.

O'Brien, Daniel P. and Steven C. Salop (2000). Competitive Effects of Partial Ownership: Financial Interest and Corporate Control. *Antitrust Law Journal* 67, 559–614.

O'Kelly, Morton E. (2012). Fuel Burn and Environmental Implications of Airline Hub Networks. *Transportation Research Part D: Transport and Environment* 17(7), 555–567.

Peltzman, Sam (2000). Prices Rise Faster Than They Fall. *Journal of Political Economy* 108(3), 466–502.

Pesaran, M. Hashem, and Ron P. Smith (1995). Estimating Long-Run Relationships from Dynamic Hetergenous Panels. *Journal of Econometrics* 69, 79–113.

Pesaran, M. Hashem, Yongcheol Shin and Ron P. Smith (1999). Pooled Mean Group Estimation of Dynamic Heterogeneous Panels. *Journal of the American Statistical Association* 94, 621–634.

Reiss, Peter, and Frank Wolak (2007). Structural Econometric Modeling: Rationales and Examples from Industrial Organisation. In *Handbook of Econometrics* Volume 6A.

Requate, Till (2006). Environmental Policy Under Imperfect Competition: A Survey. In: Tom Tietenberg and Henk Folmer (Eds.), *International Yearbook of Environmental and Resource Economics 2006/2007*, Edward Elgar, Cheltenham: UK.

Ritz, Robert A. (2014). On Welfare Losses Due to Imperfect Competition. *Journal of Industrial Economics* 62, 167–190.

Stolper, Samuel (2016). Who Bears the Burden of Energy Taxes? The Role of Local Pass-Through. Harvard Environmental Economics Program Discussion Paper 16-70.

Sutton, John (2007). Market Structure: Theory and Evidence. In: Mark Armstrong and Robert Porter (eds.), *Handbook of Industrial Organization*, Volume 3, 2301–2368.

Turner, Peter A and Siew Hoon Lim (2015). Hedging Jet Fuel Price Risk: The Case of US Passenger Airlines. *Journal of Air Transport Management* 44, 54–64.

Verboven, Frank and Theon Van Dijk (2009). Cartel Damages Claims and the Passing-on Defense. *Journal of Industrial Economics* 57, 457–491.

Vickers, John (1985). Delegation and the Theory of the Firm. *Economic Journal* 95, 138–147.

Vives, Xavier (2016). *Competition and Stability in Banking: The Role of Regulation and Competition Policy*. Princeton, NJ: Princeton University Press.

Weyl, E. Glen and Michal Fabinger (2013). Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition. *Journal of Political Economy* 121, 528–583.

Zhang, Anming and Yahua Zhang (2017). Airline Economics: An Introductory Survey. Working Paper at University of British Columbia.