

DISCUSSION PAPER SERIES

DP15127

GENERALIZED LINEAR COMPETITION: FROM PASS-THROUGH TO POLICY

Christos Genakos, Felix Grey and Robert Ritz

INDUSTRIAL ORGANIZATION



GENERALIZED LINEAR COMPETITION: FROM PASS-THROUGH TO POLICY

Christos Genakos, Felix Grey and Robert Ritz

Discussion Paper DP15127
Published 03 August 2020
Submitted 30 July 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Christos Genakos, Felix Grey and Robert Ritz

GENERALIZED LINEAR COMPETITION: FROM PASS-THROUGH TO POLICY

Abstract

Economic policy and shifts in input market prices often have significant effects on the marginal costs of firms and can prompt strategic responses that make their impact hard to predict. We introduce "generalized linear competition" (GLC), a new model that nests many existing theories of imperfect competition. We show how firm-level cost pass-through is a sufficient statistic to calculate the impact of a cost shift on an individual firm's profits. GLC sidesteps estimation of a demand system and requires no assumptions about the mode of competition, rivals' technologies and strategies, or "equilibrium". In an empirical application to the US airline market, we demonstrate GLC's usefulness for ex ante policy evaluation and identify the winners and losers of climate-change policy. We also show how GLC's structure, under additional assumptions, can be used for welfare analysis and to endogenize the extent of regulation.

JEL Classification: D43, H23, L51, L93

Keywords: Pass-Through, Imperfect Competition, regulation, Carbon Pricing, Airlines, political economy

Christos Genakos - c.genakos@jbs.cam.ac.uk
Cambridge Judge Business School, AUEB, CEP and CEPR

Felix Grey - fg313@cam.ac.uk
University of Cambridge

Robert Ritz - rar36@cam.ac.uk
University of Cambridge

Acknowledgements

Our thanks are due to Toke Aidt, Meredith Crowley, Natalia Fabra, Ken Gillingham, Par Holmberg, Matt Kotchen, Matti Liski, Hamish Low, Nathan Miller, Kamiar Mohaddes, David Newbery, Mario Pagliero, Andrea Pataconi, Andrew Rhodes, Donald Robertson, Howard Smith, Katalin Springel, Arthur van Benthem, Frank Verboven and John Vickers for helpful comments and discussions, to Severin Borenstein for providing airline data, and to Pradeep Venkatesh and Shijian Yang for excellent research assistance. We also thank seminar audiences at Cambridge, Dundee, IIOC 2019 (Boston), Loughborough, Munich, Norwich Business School, Resources for the Future, Stockholm IFN, Toulouse, 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. The usual disclaimer applies.

GENERALIZED LINEAR COMPETITION: From pass-through to policy*

CHRISTOS GENAKOS[†], FELIX GREY[‡] AND ROBERT A. RITZ[§]

July 2020

ABSTRACT. Economic policy and shifts in input market prices often have significant effects on the marginal costs of firms and can prompt strategic responses that make their impact hard to predict. We introduce “generalized linear competition” (GLC), a new model that nests many existing theories of imperfect competition. We show how firm-level cost pass-through is a sufficient statistic to calculate the impact of a cost shift on an individual firm’s profits. GLC sidesteps estimation of a demand system and requires no assumptions about the mode of competition, rivals’ technologies and strategies, or “equilibrium” . In an empirical application to the US airline market, we demonstrate GLC’s usefulness for *ex ante* policy evaluation and identify the winners and losers of climate-change policy. We also show how GLC’s structure, under additional assumptions, can be used for welfare analysis and to endogenize the extent of regulation.

JEL CLASSIFICATION: D43, H23, L51, L93.

KEYWORDS: Pass-through, imperfect competition, regulation, carbon pricing, airlines, political economy.

*Our thanks are due to Toke Aidt, Meredith Crowley, Natalia Fabra, Ken Gillingham, Par Holmberg, Matt Kotchen, Matti Liski, Hamish Low, Nathan Miller, Kamiar Mohaddes, David Newbery, Mario Pagliero, Andrea Pataconi, Andrew Rhodes, Donald Robertson, Howard Smith, Katalin Springel, Arthur van Benthem, Frank Verboven and John Vickers for helpful comments and discussions, to Severin Borenstein for providing airline data, and to Pradeep Venkatesh and Shijian Yang for excellent research assistance. We also thank seminar audiences at Cambridge, Dundee, IIOC 2019 (Boston), Loughborough, Munich, Norwich Business School, Resources for the Future, Stockholm IFN, Toulouse, 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. The usual disclaimer applies.

[†]Cambridge Judge Business School, AUEB, CEP and CEPR

[‡]Faculty of Economics and Energy Policy Research Group, University of Cambridge

[§]Corresponding author: rar36@cam.ac.uk; Cambridge Judge Business School and Energy Policy Research Group, University of Cambridge.

1. INTRODUCTION

Economic policy and shifts in input market prices often have significant effects on the marginal costs of firms and can prompt strategic responses that are hard to predict. Changing costs may be a direct objective of regulation, for example, when it puts a price on an externality such as carbon emissions or introduces a minimum wage. In other cases, like the emergence of fracking technology that has reduced energy costs for many US companies, the shift is market-driven. Such cost shifts can have important—and potentially highly heterogeneous—impacts on the profitability of firms. An *ex ante* understanding of these profit implications can be critical to the evaluation and successful implementation of policy on one hand and to the formulation of corporate strategy on the other.

However, estimating the firm-level profit impacts of a cost shift is not straightforward. In general, its profit impact will depend on the firm’s own production technology, the structure of demand, and its rivals’ strategic responses. The last factor is particularly challenging in that modelling it usually requires information on the identities of all firms, each of their production technologies, the nature of product differentiation, what variables the firms compete on, the intensity of competition in the market, and so forth. Our aim here is to present an approach that radically simplifies this problem.

In the first part of the paper, we introduce a new reduced-form model: “generalized linear competition” (GLC). We developed GLC to respond to Sutton’s (2007) call for economists to derive “predictions which are robust across a range of model specifications which are deemed reasonable.” GLC makes weaker assumptions than typical models of imperfect competition. It assumes that firm i , but not necessarily any other firm, 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. GLC also allows firm i to reduce its exposure to the cost shift: under environmental regulation, this is switching to cleaner inputs; faced with a minimum wage, it is using less labour-intensive processes.

GLC makes no assumptions about the consumer demand system, on the technologies and strategies of firm i ’s rivals, or about “equilibrium”. In this sense, our approach is consistent

with notions of bounded rationality of firms and/or consumers (Ellison, 2006; Spiegel, 2011). A further 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' choices are strategic substitutes or complements (Bulow et al., 1985).

We use GLC to quantify the impacts of a cost shift and its winners and losers. The cost shift raises firm i 's unit cost and affects those of its rivals in an arbitrary way. In the spirit of Chetty (2009), 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.¹ 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 the cost shift if and only if its pass-through is below 100%.

To see the idea underlying GLC, consider firm i which competes à la Cournot in a differentiated products market, with marginal cost $MC_i = c_i + \tau$, where τ represents a cost shifter, and a demand curve $p_i = \alpha - \beta x_i - \delta \sum_{j \neq i} x_j$. 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 the cost shifter τ tightens by $d\tau$ and this raises its rivals' marginal costs 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_j .

This basic logic extends to a rich class of oligopoly models. GLC's structure nests, among others: Cournot-Nash, Stackelberg and conjectural-variation models (with linear demand);

¹We differ from the “firm-specific” rate of pass-through used in merger analysis (Ashenfelter et al., 1998); this asks how much of the cost saving achieved only by the two merging firms is passed onto consumers. We also differ from a firm-specific pass-through rate based on a cost change incurred only by a single firm. Muehlegger and Sweeney (2020) study how pass-through estimates vary as the scope of the cost shift widens to more firms. Our GLC framework allows for arbitrary cost changes across firms.

Bertrand and Cournot models with linearly differentiated products; multi-stage models such as Allaz and Vila (1993)’s model with forward contracting; a linear version of supply function equilibrium (Klemperer and Meyer, 1989); behavioural theories of competition such as Al-Najjar et al. (2008)’s model with sunk cost bias; and models with common ownership of firms (O’Brien and Salop, 2000) which feature prominently in the current debate on the competitive impacts of institutional stock ownership (Azar et al., 2018).

GLC’s structure also applies to richer “multidimensional” models of competition. In our baseline version of GLC, firm i sells a single product into a single market at a single price. We show that its logic extends to (i) workhorse models of multiproduct Bertrand and Cournot competition and a linear version of the upgrades approach (Johnson and Myatt, 2003, 2006), (ii) multimarket competition on a network (Elliott and Galeotti, 2019), and (iii) oligopolistic price discrimination (Hazledine, 2006). Across all of these models, a multidimensional version of firm-level cost pass-through is a sufficient statistic for the profit impact.

While it is intuitive, the critical role played by pass-through is far from obvious. In recent work, Weyl and Fabinger (2013) show, in a general class of symmetric oligopoly models, that a *market-wide* rate of cost pass-through in response to a cost shift that is uniform across firms is a useful tool to understand market performance.² As Miller et al. (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 GLC family, the firm-level profit impact depends solely on firm-level pass-through—no additional information about conduct parameter(s) is needed. This simplification of incidence analysis is a primary attraction of GLC.

The second part of the paper illustrates the usefulness of the GLC framework by estimating the profit impacts of introducing carbon pricing in the domestic US airline market. This setting is important in its own right: emissions from airline travel are projected to grow well into the 21st century and economic regulation is likely as countries seek to implement internationally-agreed climate targets in a cost-effective manner. At a price of \$30 per ton of carbon dioxide, the annual “value” of US airlines’ total carbon emissions is around \$4 billion.

²Recent work by Miklos-Thal and Shaffer (2020) re-examines some results related to the concept of an exogenous increase in competition used by Weyl and Fabinger (2013).

Like many other industries, aviation is characterized by important demand, cost and conduct heterogeneities between firms. 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) which varies widely across carriers. As a result, airlines 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., low-cost vs legacy carriers) and route-level (i.e., the same carrier competes differently on different routes). This multifaceted heterogeneity makes very challenging the task of estimating the impact of new regulation.

Leveraging GLC, we estimate the firm-level profit impacts of carbon pricing in three steps. First, we estimate pass-through for individual carriers. Using quarterly ticket price data for 1,334 domestic US carrier-routes over the period 2004–2013, we estimate pass-through of fuel costs utilizing plausibly exogenous variation in fuel prices as an instrument. Our baseline specification is a standard unbalanced panel with fixed effects for each carrier-route and year-quarter; we also control for variation in demand conditions, non-fuel costs and proxies for competition. Our results show significant intra-industry heterogeneity in pass-through: the large legacy carriers (American, Delta, United and US Airways) have pass-through that is significantly smaller than 100%; by contrast, thanks to its more fuel-efficient planes, the major low-cost carrier Southwest has pass-through above 100%.

Second, we discuss and verify that GLC's assumptions are a reasonable approximation to firm i 's production technology and competitive environment in the airline setting. Third, we use GLC's sufficient-statistics results to calculate the profit impacts. At a \$30/tCO₂ carbon price, reflecting differences in pass-through, legacy carriers' profits fall by \$234 million while Southwest is a winner with a profit increase of \$98 million.³ Overall, the industry's profits decline due to climate regulation, albeit only modestly.

Finally, we show how GLC's structure can be used for welfare analysis and to endogenize the extent of regulation. A social planner sets a "political economy" carbon price in the presence of

³For comparison, the annual "value" of emissions in our sample is approximately \$400 million.

two distortions: market power (Buchanan, 1969) and political lobbying (Grossman and Helpman, 1994). We find that the resulting endogenous carbon price of \$17.68/tCO₂ lies 41% below an illustrative Pigouvian benchmark of \$30/tCO₂; under our baseline parameter values, carbon pricing achieves an overall welfare gain of \$77 million. This welfare analysis requires additional assumptions: GLC now applies to all firms in the industry and consumers are utility-maximizers; we still do not need to select a specific mode of competition.

Related literature. This paper contributes to several strands of literature. First, we introduce a different approach to the kind of structural modelling typically employed in industrial organization (Bresnahan, 1989; Berry et al., 1995; Nevo, 2001; Reiss and Wolak, 2007; Einav and Levin, 2010). By making specific assumptions about consumer demand, firms’ production technologies and the mode of competition, and then estimating a full set of primitives, structural models have been widely used for counterfactual policy analysis. This paper offers an alternative methodology for the case of cost shifts; its major advantages are allowing for model uncertainty (Sutton, 2007) and radically reducing the computational burden (Knittel and Metaxoglou, 2014).⁴ We sidestep estimation of the demand system (and the well-known challenge of determining an appropriate market definition), and show how firm-level pass-through can be sufficient information to “close the model”.⁵

Second, we add to a rich empirical literature spanning macro- and microeconomics that has estimated pass-through in response to a variety of cost shocks, including to 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 market context, it finds evidence of “incomplete” pass-through below 100% (e.g., De Loecker et al., 2016), “complete” 100% pass-through (e.g., Fabra and Reguant, 2014) as well as pass-through above 100% (e.g., Miller et al., 2017). 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

⁴Given prior knowledge of firm-level pass-through, profit impacts can be estimated with GLC using only cross-sectional data for the firms of interest. The data requirement for implementation of GLC via pass-through estimation is comparable to that for structural models: data on prices and quantities of the input factor plus an instrument for the cost shift.

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

due to differences in competition, demand and technology (Ganapati et al., 2019), our empirical results highlight *intra-industry* pass-through heterogeneity.

Third, our paper adds to a growing environmental-economics literature that studies the impacts of emissions pricing on industry. This literature has so far 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 industry-level are typically modest; this has been found using several different modeling approaches: general equilibrium (Bovenberg et al., 2005), Cournot-style oligopoly (Hepburn et al., 2013) and event study (Bushnell et al., 2013). Our GLC-based analysis extends the theory to richer modes of competition. Our empirical findings confirm modest profit impacts at the industry-level but also highlight how this masks substantial variation in the sign and extent of profit changes at the firm-level.

Fourth, we contribute to the literature on competition in the airline industry. This literature has been primarily concerned with estimating competitiveness and issues of market structure (Brander and Zhang, 1990; Kim and Singal, 1993; Berry and 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 and Syverson, 2008; Ciliberto and 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 helps understand how carriers have different incentives to influence regulation.

The rest of the paper is organized as follows. Section 2 sets out GLC, relates it to existing oligopoly models, and derives our main result on firm-level pass-through as a sufficient statistic. Section 3 extends GLC to multiple products, markets, and prices. Section 4 presents the econometric estimation of pass-through for US airlines. Section 5 discusses the empirical implementation of GLC and presents our estimates on the producer, consumer and total welfare impacts of carbon pricing. Section 6 concludes.⁶

⁶Appendix A contains proofs of our theoretical results while Appendix B contains further details on our empirical analysis. Appendix C discusses extensions to welfare analysis and endogenous regulation.

2. GENERALIZED LINEAR COMPETITION (GLC)

This section introduces a simple reduced-form model which we call “generalized linear competition” (GLC). As discussed in the introduction, we have in mind a change in policy, regulation or market prices that leads to a shift in the unit costs of firms in a particular industry. We first set out GLC’s key features and place it in the context of existing oligopoly models. We then derive our main result on the impact of a cost shift on the profits of an individual firm, explain how it compares to existing approaches in the literature, and test for its robustness using Monte Carlo analysis. For expositional clarity, the following discussion is mostly couched in terms of a cost shift that arises due to regulation.

2.1. Setup of the GLC model. Firm i competes in an industry with n firms, selling an output quantity x_i of its product at a price p_i . Let $N = \{1, 2, \dots, n\}$ denote the set of firms. 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 and Oates, 1988). Regulation then corresponds to putting a price τ on the environmental externality. Another example is minimum wage regulation for which e_i is labour input (Draca et al., 2011)

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 regulation may also apply to all (“complete regulation”), some or none of the other firms in the industry (“incomplete regulation”). More specifically, let $\phi_k \in \{0, 1\}$ be an indicator variable which equals 1 if firm k is subject to the regulation and equals 0 otherwise. Our setup has $\phi_i = 1$ for firm i but does not rely on any particular assumptions about the ϕ_j s of its rivals ($j \neq i$). Let the vector $\Phi = (\phi_k)_{k \in N}$ summarize the scope of regulation.

GLC makes four assumptions about the production technology and supply behaviour of firm i . These are taken to hold over some interval $\tau \in [\underline{\tau}, \bar{\tau}]$ of interest over which the extent of regulation varies:

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

A1 is a standard assumption which is appropriate for many forms of regulation, including a

tax on emissions.⁷

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 also a canonical assumption in microeconomic theory.⁸

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 that is common across several strands of the literature, including the theory of pass-through (Bulow and Pfleiderer, 1983; Anderson et al., 2001; Weyl and Fabinger 2013), empirical industrial organization (Berry et al., 1995; Nevo, 2001; Reiss and Wolak, 2007), environmental regulation under imperfect competition (Requate, 2006; Fowlie et al., 2016; Miller et al., 2017) and in the analysis of the profit impacts of a minimum wage (Ashenfelter and Smith, 1979; Draca et al., 2011). It rules out, at least over the range $\tau \in [\underline{\tau}, \bar{\tau}]$, the presence of (binding) capacity constraints.

Combining A1–A3, standard production theory shows that, in response to tighter regulation, $z_i(\bar{\tau}) \leq z_i(\underline{\tau})$ and $c_i(\bar{\tau}) \geq c_i(\underline{\tau})$. 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).⁹ For environmental regulation, this represents abatement: a lower emissions intensity $z_i(\bar{\tau})$ comes at a per-unit abatement cost $[c_i(\bar{\tau}) - c_i(\underline{\tau})] \geq 0$.¹⁰ For minimum-wage regulation, it represents a reduction in the labour intensity of output, achieved by substitution towards more capital-intensive processes. If such factor substitution is infeasible or unprofitable then $z_i(\bar{\tau}) = z_i(\underline{\tau})$.

⁷In Section 5, we use the GLC structure to endogenize the choice of regulation τ by government.

⁸For 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 efficiency property of market-based regulation (Baumol and Oates, 1988).

⁹We do not require any specific functional-form assumptions on the relationship between z_i and c_i .

¹⁰GLC's technology is consistent with standard properties from the environmental 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 . As usual, 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, $\partial^2 C_i(x_i, e_i) / \partial x_i \partial e_i < 0$.

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 $\bar{\tau}$, the corresponding increase in i 's optimal unit cost equals $\Delta k_i(\bar{\tau}, \underline{\tau}) = \int_{s=\underline{\tau}}^{\bar{\tau}} z_i(s) ds$.

It is useful to consider an example of GLC's A1–A3 at work. Firm i may have access to a technology that reduces its factor intensity by ζ_i at a cost of $g_i(\zeta_i)$ per unit of output. Faced with regulation τ tightening from $\underline{\tau}$ to $\bar{\tau}$ its optimized unit cost of production becomes $k_i(\bar{\tau}) = \min_{\zeta_i} [c_i(\underline{\tau}) + \bar{\tau}(z_i(\underline{\tau}) - \zeta_i) + g_i(\zeta_i)]$. Denoting the cost-minimizing value as $\zeta_i^*(\bar{\tau}) = g_i^{-1}(\bar{\tau})$, it follows that its optimal regulatory intensity becomes $z_i(\bar{\tau}) = z_i(\underline{\tau}) - \zeta_i^*(\bar{\tau})$, achieved at an incremental unit cost of $g_i(\zeta_i^*(\bar{\tau}))$.¹¹

Remark 1. *While our exposition focuses on regulation that is effectively an input tax, GLC nests an output tax as a special case where firm i 's regulatory intensity satisfies $z_i(\tau) \equiv 1$ for all $\tau \in [\underline{\tau}, \bar{\tau}]$. GLC can also apply to command-and-control regulation for which the government mandates a particular usage of inputs; an example is mandatory blending of biofuels into gasoline. 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 factor substitution is infeasible.*

Remark 2. *GLC can be applied to market-driven cost shifts such as a fall in energy costs due to fracking technology. For example, firm i may source natural gas as an input to production at a market price τ with a factor intensity $z_i(\tau)$ per unit of output. An advance in technology reduces the market price of natural gas, to which firm i responds optimally as per A1–A3. As will become clear, GLC can also accommodate firm-specific differences in input prices (that is, i 's rivals may pay different prices than τ for natural gas).*

The final assumption is the defining feature of GLC:

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.

¹¹For the cost-minimization problem to be well-behaved, we assume $g_i(0) = g_i'(0) = 0$, $g_i''(\zeta_i) > 0$, and that the optimal $g_i^{-1}(\bar{\tau}) \leq z_i(\underline{\tau})$ such that $z_i(\bar{\tau}) \geq 0$.

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. The substantive restriction here is that the proportionality factor ψ_i does not vary with the regulation τ . While regulation can shift firm i 's supply schedule it does not alter the slope of that schedule.

GLC 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). Firm i 's rivals may therefore have a degree of market power or be a competitive fringe.

GLC 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 products there are in the market (so that regulation may induce exit of firm i 's rivals and/or new entry); (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. We therefore do not have to adopt a particular market definition or impose market clearing (that is, demand equals supply at the market-level).

Hence, GLC 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 heuristics or other "rule-of-thumb" behaviour by firm i that may itself not be profit-maximizing (though we do require cost-minimization as per A2).

2.2. Special cases of GLC. To illustrate GLC's scope, we next set out examples of widely-used models of imperfect competition for which A4 is satisfied. In these models, given A1–A3 and a linear demand structure, an individual firm i 's first-order condition for profit-maximization *directly* yields a linear supply schedule that takes the form of A4.

Cournot competition with a linear market demand curve $p = \alpha - \beta \sum_{i \in N} x_i$. It is easy to check that firm i 's first-order condition directly implies $\psi_i = \beta^{-1}$ ($\forall i$). Including a firm-specific conjectural variation $v_i \equiv \left(\sum_{j \in N \setminus i} dx_j \right) / dx_i$ (Bresnahan, 1989) leads to $\psi_i = [\beta(1 + v_i)]^{-1}$, still consistent with the GLC. This also nests as a special case a 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 $p_i = \alpha_i - x_i - \delta \sum_{j \in N \setminus i} x_j$, 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 (Hackner, 2000). The latter is independent of τ as long as the set of active firms N and degree of product heterogeneity are unaffected. A richer “semi-linear” demand system $p_i = \alpha_i - \beta_i x_i - f_i(x_1, \dots, x_{j \neq i}, \dots, x_n)$ remains part of the GLC, for *any* cross-price effects implied by the function $f_i(\cdot)$.

Spatial competition on a Salop circle (Salop, 1979). A uniform mass M of consumers is distributed around a circle, where location indexes consumer preference and a linear transportation cost t serves as an index of product differentiation. There are n firms located symmetrically around the circle that compete on price. In this setting, firm i 's demand takes the linear form $D_i = M/n + (p - p_i)/t$, where its neighbours set a price equal to p , and so its first-order condition yields A4 with $\psi_i = 1/t$ ($\forall i$).

A linear version of supply function equilibrium (Klemperer and Meyer, 1989). Demand is linear $p = 1 - \sum_{i \in N} x_i$ and firm i has a linear supply schedule of the form $x_i = \sigma_i + \mu(p - k)$, where it chooses σ_i (Menezes and Quiggin, 2012). Even though the strategy space has an affine supply function, firm i 's first-order condition features A4 with $\psi_i = [1 + (n - 1)\mu]$ ($\forall i$).¹²

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 et al., 2018) and several other industries. With Cournot competition, linear demand, and assuming symmetry with $\omega_i = \omega$ ($\forall i$), this yields $\psi_i = [\beta(1 + \omega)]^{-1}$ ($\forall i$).

¹²We have modified the setup of Menezes and Quiggin (2012) slightly, by dropping their normalization in terms of production costs, without affecting any of the conclusions. The argument leading to A4 does not rely on symmetric marginal costs. If costs are symmetric, with $k_i = k$ ($\forall i$), then 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 $\mu \rightarrow 0$ and to the Bertrand paradox as $\mu \rightarrow \infty$.

GLC is more flexible than standard oligopoly models along important dimensions. First, firms within the GLC 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 quantity). Third, by allowing ψ_i to vary across firms, GLC does not impose that a firm with a higher market share necessarily has a higher profit margin. This feature is “baked into” many of the above oligopoly models via the restriction $\psi_i = \psi$ ($\forall i$).

Remark 3. *GLC can represent the outcome of tacit collusion but not necessarily the strategies supporting it. On one hand, appropriate choice of firm-level conjectural variation parameters makes it possible to generate collusive quantities and prices. On the other hand, A4 does not explicitly feature the underlying trigger strategies that might support a collusive outcome.*

Remark 4. *Despite the linearity of A1–A4, GLC does not imply that firm i ’s “pricing function” $p_i(k_1(\cdot), \dots, k_i(\cdot), \dots, k_n(\cdot))$ is necessarily linear in its arguments. For example, with Nash differentiated-products competition with linear demand $p_i = \alpha_i - x_i - \delta \sum_{j \in N \setminus i} x_j$ the pricing function is indeed linear in all firms’ marginal costs but for semi-linear demand $p_i = \alpha_i - \beta_i x_i - f_i(x_1, \dots, x_{j \neq i}, \dots, x_n)$ it is not.*

GLC also encompasses richer multi-stage models in which product-market competition occurs at the final stage of the game. Such models are commonly solved backwards to obtain the subgame-perfect Nash equilibrium. The key point is that, given A1–A3 and anticipating the choices made in subsequent stages, maximizing behaviour in the first stage yields supply schedules consistent with A4. Examples include the following:

Forward contracting and vertical integration (Allaz and Vila, 1993; Bushnell et al., 2008). In the first stage, firms can trade in a forward market; in the second stage, firms compete à la Cournot given their forward positions (Allaz and Vila, 1993). With linear demand $p = \alpha - \beta \sum_{i \in N} x_i$, the subgame-perfect equilibrium features $\psi_i = (\beta/n)^{-1}$ ($\forall i$) (Ritz, 2014); this is again independent of the regulation τ as long as n is fixed. With strategic substitutes, firms commit to selling forward a fraction of their production—which intensifies competition. A model setup in which the first stage instead involves either long-term contracts with customers or vertical integration into retail markets is strategically equivalent (Bushnell et al., 2008).

Competition with behavioural biases: misallocation of sunk costs (Al-Najjar, et al., 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. In the first stage, firms learn about the impacts of costing in a distortion game before differentiated Bertrand competition in the second stage. With linear demand $D_i = a - bp_i + g \sum_{j \in N \setminus i} p_j$ and symmetric firms, the equilibrium features $s_i = s^* > 0$ and a constant $\psi_i = \psi$ (both $\forall i$).¹³ With strategic complements, firms in equilibrium set prices that are partially based on sunk costs so as to soften competition.

The assumptions needed for such multi-stage models to form part of GLC are stronger than for the benchmark models discussed previously. The reason is that firm i 's supply schedule taking the form of A4 hinges on Nash-maximizing behaviour of i 's rivals in the second stage; this determines their aggregate best response in the product market which, in turn, is what firms in the first stage seek to strategically influence. By contrast, in simpler oligopoly models, profit-maximization by firm i alone yields a first-order condition that directly takes the form of A4—irrespective of the behaviour of i 's rivals.

Remark 5. *GLC is conceptually distinct from classes of oligopoly models that are aggregative games (Corchón, 1994; Acemoglu and Jensen, 2013) or potential games (Monderer and Shapley, 1996). To see this, note that Cournot-Nash competition with linear market demand $p = \alpha - \beta \sum_{i \in N} x_i$ is an aggregative game, a potential game, and also a member of GLC. However, with differentiated-products demand for firm i of $p_i = \alpha - \beta_i x_i - \sum_{j \in N \setminus i} \delta_{ij} x_j$, it is no longer aggregative nor a potential game (unless $\delta_{ij} = \delta_{ji}$ for all i and $j \neq i$) but still yields A4.*

Remark 6. *GLC is consistent with the standard paradigm that firms base their decision-making on marginal costs. However, GLC does not require this assumption and instead can allow for alternatives such as average-cost pricing. Simply think of firm i 's unit cost $k_i(\tau) = c_i(\tau) + \tau z_i(\tau)$ from A3 instead as its average cost where $c_i(\tau)$ also incorporates its fixed costs. All else equal, regulation τ raises firm i 's marginal and average cost by the same amount. A1 applies in the same way, A2 then requires that firm i 's use of technology minimizes its average cost, and A4 then says that i 's product-market strategy is based on average cost.*

¹³Specifically, using the results in Section 5 of Al-Najjar et al. (2008), $\psi_i = [b/(2b + g)] [b(1 + g/[2b - (n - 1)g]) - (n - 1)g^2/[2b - (n - 1)g]]$.

2.3. The profit impact of a cost shift. We wish to quantify the impact of regulation τ on firm i 's profits. Suppose that the extent of regulation rises from initial level $\underline{\tau}$ to a higher $\bar{\tau}$. 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, $\bar{\tau} \rightarrow \underline{\tau}$.

Let $\Pi_i(\tau; \Phi)$ denote firm i 's optimized profits as a function of regulation, and similarly $e_i(\tau)$ is its optimal use of the regulated factor, and define $\Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \equiv [\Pi_i(\bar{\tau}; \Phi) - \Pi_i(\underline{\tau}; \Phi)]$ as the change in profits due to regulation. Observe that, if firm i does not respond to the tightening of regulation in any way, and its rivals do not change their behaviour either, then its profits simply decline according to $\Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) = -[(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})]$, that is, by the “static” impact of regulation associated with its *initial* quantity of the regulated factor. This static benchmark is easy to calculate insofar as minimal information is available on firm i 's initial position. Our main objective is to find a parsimonious result for the “dynamic” profit impact that takes into account the responses of firm i and its rivals.¹⁴

We will see that firm-level pass-through plays a central role for the profit impact of regulation. Define firm i 's *marginal* rate of cost pass-through as:

$$\rho_i(\tau; \Phi) \equiv \frac{dp_i(\tau; \Phi)/d\tau}{dk_i(\tau)/d\tau}. \quad (1)$$

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 the cost increases incurred by firm i and its rivals, as reflected by the scope of regulation $\Phi = (\phi_k)_{k \in N}$, as well as by their demand conditions and product-market behaviour.

In standard oligopoly models, in which A1–A4 apply to all n firms, given the scale τ and scope Φ of regulation, firm i 's equilibrium price $p_i(\tau; \Phi) = p_i(k_1(\phi_1\tau), \dots, k_i(\tau), \dots, k_n(\phi_n\tau))$ is a function of the marginal costs of all firms. So the price response $dp_i(\tau; \Phi)/d\tau$ captures any relevant changes in these costs, also reflecting the scope Φ of regulation. Our firm-level pass-through rate therefore reflects the impact of regulation on all players, including heterogeneity

¹⁴Our use of the “static” vs “dynamic” terminology here is in the spirit of the distinction between static vs dynamic scoring in the public finance literature, where the latter takes into account behavioural responses in response to a tax cut rather than treating the tax base as fixed (see, e.g., Mankiw and Weizierl, 2006). The GLC model setup is static in a temporal sense.

across firms. It is distinct from the market-wide pass-through rate considered in much of the literature for which firms' marginal costs rise uniformly.¹⁵ For a given price change, a firm that experiences half the cost increase of a rival has firm-level pass-through twice as high.

Two further definitions will be useful. First, denote i 's *average* rate of cost pass-through over the interval $\tau \in [\underline{\tau}, \bar{\tau}]$ as:

$$\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) \equiv \frac{\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi)}{\Delta k_i(\bar{\tau}, \underline{\tau})} = \frac{\int_{\tau=\underline{\tau}}^{\bar{\tau}} [\rho_i(\tau; \Phi) z_i(\tau)] d\tau}{\int_{\tau=\underline{\tau}}^{\bar{\tau}} z_i(\tau) d\tau} \quad (2)$$

Average pass-through corresponds to the discrete changes in prices and costs that are typically observed in empirical data. For small changes in regulation $\bar{\tau} \rightarrow \underline{\tau}$, average and marginal cost pass-through are locally approximately equal, $\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) \simeq \rho_i(\underline{\tau}; \Phi)$. Second, define an inverse measure of the extent of factor substitution over the interval $\tau \in [\underline{\tau}, \bar{\tau}]$ as $g_i(\bar{\tau}, \underline{\tau}) \equiv [\int_{\tau=\underline{\tau}}^{\bar{\tau}} z_i(\tau) d\tau] / [(\bar{\tau} - \underline{\tau}) z_i(\underline{\tau})] > 0$. We have $g_i(\bar{\tau}, \underline{\tau}) = 1$ for all $\tau \in [\underline{\tau}, \bar{\tau}]$ if regulation is an output tax (for which $z_i(\tau) \equiv 1$ for all $\tau \in [\underline{\tau}, \bar{\tau}]$) or if factor substitution is infeasible. For a small tightening of regulation, $\bar{\tau} \rightarrow \underline{\tau}$, we have $g_i(\bar{\tau}, \underline{\tau}) \simeq 1$.

We thus obtain our first main result for GLC.

Proposition 1. Under GLC, as defined by A1–A4, regulation τ with scope Φ affects firm i 's profits Π_i according to:

(a) For a “small” tightening of regulation from $\underline{\tau}$ to $\bar{\tau}$ (with $\bar{\tau} \rightarrow \underline{\tau}$):

$$\Delta \Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \Big|_{\bar{\tau} \rightarrow \underline{\tau}} \simeq (\bar{\tau} - \underline{\tau}) \frac{d\Pi_i(\tau; \Phi)}{d\tau} \Big|_{\tau=\underline{\tau}} = -2[1 - \rho_i(\underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau}) e_i(\underline{\tau})].$$

(b) In general, for a “large” tightening of regulation from $\underline{\tau}$ to $\bar{\tau}$:

$$\Delta \Pi_i(\bar{\tau}, \underline{\tau}; \Phi) = -2 \int_{\tau=\underline{\tau}}^{\bar{\tau}} [1 - \rho_i(\tau)] e_i(\tau) d\tau = -2[1 - \bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi)] \Omega_i(\bar{\tau}, \underline{\tau}; \Phi) [(\bar{\tau} - \underline{\tau}) e_i(\underline{\tau})],$$

where:

$$\Omega_i(\bar{\tau}, \underline{\tau}; \Phi) \equiv g_i(\bar{\tau}, \underline{\tau}) \left(1 - \frac{g_i(\bar{\tau}, \underline{\tau}) [1 - \bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi)] [(\bar{\tau} - \underline{\tau}) z_i(\underline{\tau})] / p_i(\underline{\tau})}{2[p_i(\underline{\tau}) - k_i(\underline{\tau})] / p_i(\underline{\tau})} \right) > 0.$$

(c) Suppose that firm i 's (i) cost increase is modest relative to its initial price, that is, $[(\bar{\tau} -$

¹⁵If 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.

$\tau)z_i(\tau)]/p_i(\tau)$ is “small”, and (ii) regulatory intensity is approximately constant, that is, $z_i(\bar{\tau}) \simeq z_i(\tau)$. Then, for a “large” tightening of regulation from τ to $\bar{\tau}$:

$$\Delta\Pi_i(\bar{\tau}, \tau; \Phi) \simeq -2[1 - \bar{\rho}_i(\bar{\tau}, \tau; \Phi)][(\bar{\tau} - \tau)e_i(\tau)].$$

Part (a) of Proposition 1 makes precise how, to first order, firm i 's rate of cost pass-through *alone* is a sufficient statistic for the profit impact of regulation τ . That is, firm-level pass-through is the only thing that matters for going from the “static” estimate of the profit impact (i.e., $\Delta\Pi_i(\bar{\tau}, \tau; \Phi) = -[(\bar{\tau} - \tau)e_i(\tau)]$) to the “dynamic” GLC result. The resulting very simple expression for the profit impact applies to all models that are part of the GLC family.

Firm i 's pass-through rate captures all relevant information about supply and demand: 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. In our setting, pass-through therefore also captures the import of behavioural biases or “irrationality” on the part of firms and/or consumers. Whatever their other differences, if any two theories within the GLC imply identical pass-through for firm i , then they also imply an identical profit impact (for a given initial $e_i(\tau)$).

To understand the result, recall that 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 i 's pass-through rate captures. These two roles drive the “twoness” of $\Delta\Pi_i(\bar{\tau}, \tau; \Phi) \simeq -2[1 - \rho_i(\tau; \Phi)][(\bar{\tau} - \tau)e_i(\tau)]$. Therefore pass-through *signs* the profit impact: i 's profits fall if and only if its pass-through rate is less than 100%; then its profit margin shrinks and, by A4, it also experiences weaker sales. Conversely, with pass-through above 100%, the firm benefits from tighter regulation as both its profit margin and sales volume rise.

While it is intuitive, this critical role played by pass-through is also far from obvious. Weyl and 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

and Donaldson, 2015). The profit impact that corresponds to Proposition 1(a) 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 a (symmetrized) price elasticity of market demand. By contrast, within GLC, the first-order 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 GLC. Compared to existing literature, GLC allows for near-arbitrary heterogeneity across firms but makes heavy use of the linear structure implied by A4.

Another feature of the result is that the first-order profit impact $\Delta\Pi_i$ is independent of the proportionality term ψ_i from A4. As we have seen, different models within the GLC family differ in terms of their implied ψ_i . Yet Proposition 1(a) 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 $dx_i(\tau) = \psi_i [dp_i(\tau) - dk_i(\tau)]$ 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 initial use of the regulatory factor $e_i(\underline{\tau})$ does not depend on ψ_i . This is also one reason for why Proposition 1(a) applies without requiring separate information about own-price and cross-price elasticities of demand.¹⁶

Part (b) provides a general result that applies for a “large” tightening of regulation. The profit impact can be written as the integral of marginal profit impacts over the interval $\tau \in [\underline{\tau}, \bar{\tau}]$, starting from the “small” tightening characterized by part (a). While instructive, the resulting integral is less useful because it does not immediately relate to firm i ’s initial use of regulated factor $e_i(\underline{\tau})$. The next expression therefore restates the profit impact relative to the familiar static estimate (i.e., $-(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})$); the dynamic effects are now captured by the term $2[1 - \bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi)]\Omega_i(\bar{\tau}, \underline{\tau}; \Phi)$. Analogous to before, it is now the average pass-through rate, relative to 100%, that signs the profit impact. The additional term $\Omega_i(\bar{\tau}, \underline{\tau}; \Phi)$ depends on several other firm characteristics such as the factor-substitution measure $g_i(\bar{\tau}, \underline{\tau})$ and firm i ’s initial price-cost markup $[p_i(\underline{\tau}) - k_i(\underline{\tau})]/p_i(\underline{\tau})$.

¹⁶Of 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(a), 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).

As emphasized by Chetty (2009), surplus measures in economic models are typically highly non-linear functions; as a consequence, sufficient-statistics approaches like ours tend to focus on first-order surplus impacts like in part (a)—which can be seen as the outcome of linearization around an initial equilibrium. The result in part (b) begins to morph back into a structural approach, notably in that it includes terms such as firm i 's price-cost markup. Nonetheless, Proposition 1(b) still captures GLC's spirit of simplifying and highlighting the critical role played by pass-through at the firm-level.

Part (c) provides a “mid-way” result with a particular view towards empirical implementation. The basic challenge at this juncture is that, on one hand, the formula for the profit impact from part (a) is highly appealing for its simplicity while, on the other hand, the more involved result from part (b) may be more relevant insofar as cost shifts observed in empirical data are not literally infinitesimal. Part (c) bridges these two results by making two additional assumptions. First, a “static” measure of the size of i 's (non-infinitesimal) cost shift, *relative* to its initial product price, $[(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})]/p_i(\underline{\tau})$ should be “small”. This should be relatively easy to check in many applications, based on “pre-shock” data. Second, either the design of regulation or firm i 's production technology should imply that its regulatory intensity is approximately constant over the interval, $z_i(\bar{\tau}) \simeq z_i(\underline{\tau})$, so that $g_i(\bar{\tau}, \underline{\tau}) \simeq 1$. This may follow from the institutional context or may be checked in the data. Taken together, these assumptions imply that $\Omega_i(\bar{\tau}, \underline{\tau}; \Phi) \simeq 1$, and so the estimated profit impact $\Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \simeq -2[1 - \bar{p}_i(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})]$ again takes the same simple form as in part (a). The validity of this approach will depend on the particular application being considered but we hope that it will prove particularly useful in practice.

Remark 7. *GLC accommodates other margins of adjustment that can be challenging for empirical work. First, consumer preferences may vary as a by-product of regulation; for example, following the introduction of a soda tax, some consumers may discover that healthier drinks are not so bad after all—and therefore alter their purchasing behaviour. Second, regulation may induce firm i 's rivals to reposition their products; for example, faced with climate regulation, firms may redesign products to be “greener”. In both cases, Proposition 1 can continue to apply. To illustrate the logic, write firm i 's demand as $p_i = \alpha_i(\tau) - \beta_i x_i - \delta_i(\tau) \sum_{j \in N \setminus i} x_j$ (with $\delta_i \leq \beta_i$) so regulation may indirectly affect firm i 's demand conditions by changing willingness-to-pay $\alpha_i(\tau)$*

and/or the differentiation parameter $\delta_i(\tau)$. The key point is that, firm i 's first-order condition for quantity still satisfies A4 so pass-through remains a (first-order) sufficient statistic.

Up to this point, we have treated i 's pass-through rate as a parameter. To progress further, there are two basic approaches. The first—albeit against the spirit of GLC—is to select a specific model of competition and derive the theory-based rate of pass-through. The second is to combine the structural result from Proposition 1 with firm-level empirical estimates of pass-through; we pursue this approach later on for the US airline industry.

2.4. Robustness to non-linearities: Monte Carlo results. The GLC structure relies heavily on linearity of firm i 's costs and its supply-side behaviour. We now use Monte Carlo simulations to explore quantitatively its robustness to model misspecification. The setting features two departures from GLC's assumptions: firms may have non-constant marginal cost (A3 fails) and the demand system may be non-linear (A4 fails). The idea is that a researcher can estimate pass-through but does not know the true model of competition so relies on GLC.

Following Hepburn et al. (2013), we assume that the true model is an augmented version of homogenous-product Cournot competition. The salient additional model parameters are as follows. The inverse demand curve is $p(X)$ where X is total industry output; let $\xi \equiv -Xp''(X)/p'(X)$ be an index of demand curvature, so demand is convex (concave) if $\xi \geq 0$ ($\xi < 0$), and firm i 's market share is given by $s_i \equiv x_i/X \in (0, 1)$. There is a competitiveness parameter $\theta > 0$ which nests Cournot-Nash behaviour when $\theta = 1$ and lower values of θ correspond to more intense competition. Finally, the slope of the marginal cost function is given by m and is identical across firms, and $\bar{m} \equiv -m/p' \geq 0$ is a measure of cost convexity.

For simplicity, we focus on the robustness of our first-order result from Proposition 1(a) for a marginal change in “complete” regulation that covers all firms in the market (i.e., $\Phi \equiv 1$). Define a profit-impact factor $\gamma_i \equiv 2(1 - \rho_i)$ that captures the wedge between the “static” and “dynamic” profit impacts such that $\Delta\Pi_i(\bar{\tau}, \underline{\tau}) \simeq -\gamma_i[(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})]$ under GLC. Translating Proposition 7 of Hepburn et al. (2013) into our context yields firm i 's true profit-impact parameter $\tilde{\gamma}_i$ in the presence of demand and cost non-linearities:

$$\tilde{\gamma}_i(\bar{\tau}, \underline{\tau})|_{\bar{\tau} \rightarrow \underline{\tau}} = 2(1 - \rho_i) + \frac{\theta}{(\theta + \bar{m})} \theta s_i \xi \rho_i - \frac{\bar{m}}{(\theta + \bar{m})} (1 - \rho_i),$$

where Cournot equilibrium pass-through is given by $\rho_i = [n/(n + \theta(1 - \xi) + \bar{m})] \left(\frac{1}{n} \sum_{j=1}^n z_j \right) / z_i$ and $(\theta + \bar{m}) > 0$ by stability of equilibrium. Note that pass-through incorporates the effects of both demand curvature (Bulow and Pfleiderer, 1983; Seade, 1985) and cost curvature. With linear demand and constant marginal costs ($\xi = \bar{m} = 0$), A3 is met and each firm's first-order condition satisfies A4 so Proposition 1(a) applies: $\tilde{\gamma}_i = 2(1 - \rho_i) \equiv \gamma_i$. The first additional effect stems from demand: if true demand is convex with $\xi > 0$ then this pushes $\tilde{\gamma}_i$ up so GLC's γ_i underestimates the adverse profit shock. This effect is quantitatively modest if either the market is very competitive (low θ) or firm i 's market share s_i is small or the pass-through rate ρ_i itself is small. The second additional effect arises from costs: if these are convex with $\bar{m} > 0$ this pushes down $\tilde{\gamma}_i$ (as long as $\rho_i \leq 1$) so GLC then yields an overestimate.

Two qualitative conclusions follow. First, GLC exhibits no systematic bias: depending on the precise way in which A3 and/or A4 are violated, it may over- or underestimate the true profit impact. Second, the impacts of plausible departures from A3 may partially offset those of A4. Most demand curves used by economists are (weakly) convex and it is probably true that the case of convex costs, at equilibrium, is more likely than concave costs (especially for emissions-intensive industries). These two departures from GLC's assumptions tend to work in opposite directions such that $\tilde{\gamma}_i \approx 2(1 - \rho_i) \equiv \gamma_i$ may still be a reasonable approximation.

We now use Monte Carlo analysis to further explore misspecification in a quantitative manner. In the following design, we create 10,000 hypothetical industries for which, on average, A3 is violated for each firm and so A4 is also violated for each firm. The design also yields significant heterogeneity in firms' marginal costs and regulatory intensities.¹⁷ In particular, for each industry, we draw six sets of parameters to reflect model uncertainty underlying GLC:

1. The number of firms is drawn uniformly as $n \in [2, 8]$, with an integer constraint, to create a relatively concentrated market structure;
2. The market competitiveness parameter is drawn uniformly as $\theta \in [0, 1]$ so that competition, on average, lies mid-way between perfect competition and Cournot-Nash;
3. For each of the $n \in [2, 8]$ firms in the industry, firm i 's market share is drawn uniformly

¹⁷Second-order conditions and stability conditions are always satisfied given our parameter assumptions below.

as $s_i \in (0, 1)$; then firms' market shares are re-normalized so that they sum to 100%;

4. The cost convexity metric is drawn uniformly as $\bar{m} \in [0, 1]$ leading to modestly convex costs so that A3 is violated;
5. Demand curvature ξ is drawn uniformly from five discrete values $\{-1, -\frac{1}{2}, 0, \frac{1}{2}, 1\}$, in line with the common assumption from economic theory that demand is log-concave $\xi \leq 1$ (Bagnoli and Bergstrom, 2005). This range includes three convex demand systems used in an influential paper on oligopolistic pricing by Genesove and Mullin (1998): linear ($\xi = 0$), quadratic ($\xi = \frac{1}{2}$), exponential ($\xi = 1$) as well as their concave counterparts. While linear demand ($\xi = 0$) remains focal, A4 is nonetheless violated as A3 is already violated;
6. For each of the $n \in [2, 8]$ firms in the industry, firm i 's relative regulatory intensity $\bar{z}_i \equiv z_i / \left(\frac{1}{n} \sum_{j=1}^n z_j\right)$ is drawn uniformly as $\bar{z}_i \in [\frac{1}{2}, \frac{3}{2}]$ so that each firm experiences a cost shock between 50% smaller or larger than the industry average.

We then draw one firm from each industry, such that we have 10,000 firms, and complete the analysis of the true model and the misspecified GLC in four steps. For each firm, we calculate: (1) the true firm-level pass-through rate $\rho_i(n, \theta, \xi, \bar{z}_i)$ ("true CPT"), (2) the true profit-impact factor $\tilde{\gamma}_i(\rho_i, \theta, s_i, \xi, \bar{m})$ ("true PIF"), (3) GLC's estimated profit-impact factor $\gamma_i(\rho_i) = 2(1 - \rho_i)$ ("GLC PIF"), and (4) GLC's error $\varepsilon_i \equiv \gamma_i - \tilde{\gamma}_i$. A crucial point is that GLC's PIF picks up the variation in the true CPT even if A3 and A4 are violated.¹⁸

Table 1 summarizes parameter draws and results. Firm-level pass-through rates ρ_i vary widely between around 30% up to almost 200%, partly driven by variation in the size of a firm's cost shock relative to its rivals. Similarly, there is large variation in true PIFs, with some firms being hit hard ($\tilde{\gamma}_i > 1$, so the dynamic profit impact is worse than the static impact) and others benefitting from regulation ($\tilde{\gamma}_i < 0$). Overall, the average GLC PIF of $\gamma_i = .205$ lies close to the average true PIF of $\tilde{\gamma}_i = .156$. Figure 1 is a kernel density plot of the true PIF and GLC PIF over all draws: the means are strikingly close. The true PIF has a slightly larger spread but overall distributions are quite similar in shape.

¹⁸See Miller et al. (2016) for an analysis of noisy or biased pass-through estimation.

In sum, this analysis suggests that, as long as a researcher is able to obtain good estimates of cost pass-through, GLC can perform reasonably well even in situations with significant model misspecification on the demand and/or cost side.

Remark 8. *In a differentiated-products context, the empirical industrial-organization literature (Berry et al., 1995; Nevo, 2001; Reiss and Wolak, 2007) employs discrete-choice models with a logit demand structure. This literature typically makes assumptions equivalent to GLC’s A1–A3: firms are cost-minimizers with constant marginal costs of production that take input prices as given. However, the non-linearity of the logit demand leads to a firm’s first-order condition for price departing from A4. The literature on merger analysis has used Monte Carlo simulations to explore functional-form assumptions on demand: Crooke et al. (1999) and Miller et al. (2016) both find that the estimated consumer price changes due to a merger are roughly equal under linear and logit demand. This suggests that GLC’s structure does not necessarily yield very different conclusions than existing empirical IO approaches. An advantage of GLC is its semi-parametric nature: unlike fully-specified structural models, Proposition 1 requires functional-form assumptions to be made only for firm i .*

3. MULTIDIMENSIONAL GLC

In our baseline model of GLC, firm i sells a single product into a single market at a single price. Our “multidimensional” extension to GLC in this section shows how the basic insight—firm-level pass-through as a sufficient statistic for the profit impact of a cost shift—applies in settings with multiple products, multiple markets and multiple prices. While GLC, of course, cannot be a fully general model of competition—given its reliance on supply-side linearity—these further results considerably extend its reach.

3.1. Model setup. Firm i competes in an industry with the scope of its offering captured by the set M . For each “component” $m \in M$ of its offering, it sells an output quantity x_{im} at a price p_{im} . The scope of its offering can capture, for example: (1) the same product being offered in different markets (multimarket competition); (2) the same product being offered at different prices in the same market (price discrimination); or (3) a range of products being offered in the same market (multiproduct competition).

Generalizing baseline GLC, let e_{im} be one of the inputs that i uses in production of element $m \in M$ so its total use of this factor of production is $e_i = \sum_{m \in M} e_{im}$. Regulation imposes a cost τ on each unit of this input e_i . There a set of N_m firms that are active in component m of its offering (including firm i itself). Let $\phi_{km} \in \{0, 1\}$ be an indicator variable which equals 1 if firm $k \in N_m$ that is active in firm i 's component m is subject to the regulation and equals 0 otherwise. Let $\Phi = (\phi_{km})_{k \in N_m, m \in M}$ summarize the scope of regulation across firm i 's rivals. With complete regulation, $\phi_{km} = 1$ for all $k \in N_m, m \in M$.

In general, firm i 's profits across its offering are $\Pi_i \equiv \sum_{m \in M} \Pi_{im} = \sum_{m \in M} p_{im}x_{im} - C_i(\mathbf{x}_i, \mathbf{e}_i) - \tau e_i$, where $p_{im}x_{im}$ is its sales revenue on component $m \in M$, and its total costs are made up of its operating costs $C_i(\mathbf{x}_i, \mathbf{e}_i)$, where $\mathbf{x}_i = (x_{im})_{m \in M}$, $\mathbf{e}_i = (e_{im})_{m \in M}$ are the vectors of its output and factor use, plus its regulatory costs τe_i .

Multidimensional GLC extends A1–A4 from baseline GLC about the production technology and supply behaviour of firm i across each component $m \in M$ of its offering, again taken to hold over some interval $\tau \in [\underline{\tau}, \bar{\tau}]$ of interest:

A1M. (*Input price-taking*) Firm i takes input prices, including the regulation τ , as given.

A2M. (*Cost-minimizing inputs*) Firm i chooses its inputs, including its use of the regulated factor e_i , optimally so as to minimize its total costs $C_i(\mathbf{x}_i, \mathbf{e}_i) + \tau e_i$ of producing its output vector \mathbf{x}_i .

A3M. (*Constant returns to scale*) Firm i 's optimized total costs faced with regulation τ are linear in outputs $C_i(\mathbf{x}_i, \mathbf{e}_i) + \tau e_i = \sum_{m \in M} k_{im}(\tau)x_{im}$, with the unit cost of component m given by $k_{im}(\tau) = c_{im}(\tau) + \tau z_{im}(\tau)$ where $c_{im}(\tau)$ is its per-unit operating cost and $z_{im} \equiv e_{im}/x_{im}$ is its regulatory intensity (use of the regulated factor per unit of output)

A4M. (*Linear product market behaviour*) Firm i 's product market behaviour regarding component m satisfies $x_{im}(\tau) = \psi_{im} [p_{im}(\tau) - k_{im}(\tau)]$, where $\psi_{im} > 0$ is a constant and $[p_{im}(\tau) - k_{im}(\tau)] > 0$ is its profit margin.

The economics of these four assumptions is analogous to A1–A4 in baseline GLC. A1M and A2M are standard: firm i is a cost-minimizer and a price-taker in input markets. A3M extends the notion of constant marginal costs to a multidimensional setting, allowing for the regulatory intensity z_{im} to vary across the components of i 's offering. For example, its products may have

different emissions intensities—and hence different exposures to an emissions tax. Finally, A4M is again the defining supply-linearity feature of GLC; the parameter ψ_{im} may also vary across i 's offering. Relative to baseline GLC, an additional (implicit) assumption here is that the scope of firm i 's offering M does not vary with regulation τ .

Multidimensional GLC also makes no assumptions about consumer preferences and the demand system, rivals' technology and behaviour, (strategic) substitutes vs (strategic) complements, the equilibrium concept, and so on. We show below that, given A1M–A3M, the core assumption of A4M is satisfied for a range of models from the existing industrial-organization literature that have multidimensional features—including recent models in the literature of multiproduct competition, price discrimination, and networked markets.

3.2. Main result. We wish to quantify the impact of regulation τ on firm i 's profits under multidimensional GLC, and a multidimensional version of firm-level cost pass-through will play a central role. Let $\Pi_i(\tau)$ denote firm i 's optimized overall profits as a function of regulation, and similarly $e_i(\tau)$ is its optimal use of the regulated factor.

We now introduce three sets of useful definitions. First, define firm i 's marginal rate of cost pass-through in component $m \in M$ of its offering as:

$$\rho_{im}(\tau; \Phi) \equiv \frac{dp_{im}(\tau; \Phi)/d\tau}{dk_{im}(\tau)/d\tau} \quad (3)$$

The denominator captures by how much i 's optimal unit cost $k_{im}(\tau)$ responds to a small tightening in regulation; given A1M–A3M, the envelope theorem implies $dk_{im}(\tau)/d\tau = z_{im}(\tau)$. The numerator captures the change in i 's product price on component m . Denote i 's average rate of cost pass-through on m over the interval $\tau \in [\underline{\tau}, \bar{\tau}]$ as:

$$\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) \equiv \frac{\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi)}{\Delta k_{im}(\bar{\tau}, \underline{\tau})} = \frac{\int_{\tau=\underline{\tau}}^{\bar{\tau}} [\rho_{im}(\tau; \Phi) z_{im}(\tau)] d\tau}{\int_{\tau=\underline{\tau}}^{\bar{\tau}} z_{im}(\tau) d\tau} \quad (4)$$

For small changes in regulation $\bar{\tau} \rightarrow \underline{\tau}$, average and marginal cost pass-through are locally approximately equal, $\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) \simeq \rho_{im}(\underline{\tau}; \Phi)$.

Second, we translate these pass-through definitions into a multidimensional context. Define firm i 's multidimensional rate of marginal cost pass-through as a weighted average of its pass-

through across the components of its offering:

$$\rho_i^M(\tau; \Phi) \equiv \sum_{m \in M} \omega_{im}(\tau) \rho_{im}(\tau; \Phi) \quad (5)$$

where the weights are by its use of the regulated factor, $\omega_{im}(\tau) \equiv e_{im}(\tau)/e_i(\tau) \in (0, 1)$ so that $\sum_{m \in M} \omega_{im}(\tau) = 1$. We also define a multidimensional rate of average cost pass-through as follows:

$$\bar{\rho}_i^M(\bar{\tau}, \underline{\tau}; \Phi) \equiv \sum_{m \in M} \omega_{im}(\underline{\tau}) \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi). \quad (6)$$

Note that the weights $\omega_{im}(\underline{\tau})$ here are by firm i 's initial use of the regulated factor. This will be the relevant definition for our next result and has the advantage of typically being observable “pre-shock” from an empirical perspective.

Third, define an inverse measure of the extent of factor substitution on component m over the interval $\tau \in [\underline{\tau}, \bar{\tau}]$ as $g_{im}(\bar{\tau}, \underline{\tau}) = [\int_{\tau=\underline{\tau}}^{\bar{\tau}} z_{im}(\tau) d\tau] / [(\bar{\tau} - \underline{\tau}) z_{im}(\underline{\tau})] > 0$. We have $g_{im}(\bar{\tau}, \underline{\tau}) = 1$ for all $\tau \in [\underline{\tau}, \bar{\tau}]$ and $m \in M$ if regulation is an output tax (for which $z_{im}(\tau) \equiv 1$ for all $\tau \in [\underline{\tau}, \bar{\tau}]$ and $m \in M$) or if factor substitution is infeasible. For a small tightening of regulation, $\bar{\tau} \rightarrow \underline{\tau}$, we have $g_{im}(\bar{\tau}, \underline{\tau}) \simeq 1$ for all $m \in M$.

This delivers a direct analog to Proposition 1 from baseline GLC:

Proposition 2. Under multidimensional GLC, as defined by A1M–A4M, a regulation τ with scope Φ affects firm i 's profits Π_i according to:

(a) For a “small” tightening of regulation from $\underline{\tau}$ to $\bar{\tau}$ (with $\bar{\tau} \rightarrow \underline{\tau}$):

$$\Delta \Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \Big|_{\bar{\tau} \rightarrow \underline{\tau}} \simeq (\bar{\tau} - \underline{\tau}) \frac{d\Pi_i(\tau; \Phi)}{d\tau} \Big|_{\tau=\underline{\tau}} = -2[1 - \rho_i^M(\underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})].$$

(b) In general, for a “large” tightening of regulation from $\underline{\tau}$ to $\bar{\tau}$:

$$\Delta \Pi_i(\bar{\tau}, \underline{\tau}; \Phi) = -2 \sum_{m \in M} [1 - \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})e_{im}(\underline{\tau})]\Omega_{im}(\bar{\tau}, \underline{\tau}; \Phi),$$

where:

$$\Omega_{im}(\bar{\tau}, \underline{\tau}; \Phi) \equiv g_{im}(\bar{\tau}, \underline{\tau}) \left(1 - \frac{g_{im}(\bar{\tau}, \underline{\tau})[1 - \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})z_{im}(\underline{\tau})/p_{im}(\underline{\tau})]}{2[p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})]/p_{im}(\underline{\tau})} \right) > 0.$$

(c) Suppose that for each component $m \in M$ of firm i 's offering (i) its cost increase is modest

relative to its initial price, that is, $[(\bar{\tau} - \underline{\tau})z_{im}(\underline{\tau})]/p_{im}(\underline{\tau})$ is “small”, and (ii) regulatory intensity is approximately constant, that is, $z_{im}(\bar{\tau}) \simeq z_{im}(\underline{\tau})$. Then, for a “large” tightening of regulation from $\underline{\tau}$ to $\bar{\tau}$:

$$\Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \simeq -2[1 - \bar{\rho}_i^M(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})].$$

Proposition 2 establishes firm-level cost pass-through as a sufficient statistic in multidimensional settings. As under baseline GLC, the “static” profit impact is given by $\Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) = -(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})$, where firm i ’s initial use of the regulated factor covers all components of its offering. In direct parallel to Proposition 1, Proposition 2’s parts (a)-(c) present three ways of going to the “dynamic” version that captures the strategic responses by firm i and its rivals.

Part (a) gives a first-order approximation that shows how i ’s multidimensional (marginal) pass-through rate determines its profit impact. The intuition is as under baseline GLC: given the linearity of i ’s costs and supply, firm-level pass-through captures both margin and sales impacts for each individual component—and the weights then aggregate up to obtain a multidimensional profit impact. Once again, therefore, pass-through signs the profit impact: i ’s profits fall if and only if its multidimensional pass-through rate is less than 100%.

Under multidimensional GLC, there is an additional “portfolio effect” that is not present in the baseline GLC. The profit impact for firm i is relatively more favourable to the extent that more important offerings (with higher weights ω_{im}) also have higher pass-through (higher ρ_{im}), as this raises its multidimensional pass-through ρ_i^M which in turn raises $\Delta\Pi_i$ (for a given overall exposure to regulation, e_i).

Part (b) presents a general result that applies to discrete changes in regulation.

Part (c), with an eye towards empirical implementation, introduces two additional conditions that are straightforward generalizations of those in Proposition 1(c): for each component $m \in M$ of firm i ’s offering, the cost shocks due to regulation are “small” relative to product prices and there is limited scope for factor substitution. Under these conditions, the simplicity of the result in part (a) is restored and multidimensional (average) pass-through alone re-emerges as a sufficient statistic. Empirical implementation therefore does not require knowledge of firm i ’s supply parameters, $(\psi_{im})_{m \in M}$ from A4M—or of conduct parameters and price and cross-price elasticities of demand. Note also that the weights in part (a) and part (c) both reflect the initial

exposure $\omega_{im}(\tau) = e_{im}(\tau)/e_i(\tau)$ of each component.

It is worth emphasizing that A1M–A4M do allow for interdependencies between the M components of i 's offering; put differently, Proposition 2 is not a trivial extension to Proposition 1. To see why, consider a duopoly selling two products A and B . In a standard model of multiproduct pricing, firm i 's equilibrium price for product A can be written as $p_i(k_{iA}(\tau), k_{jA}(\tau\phi_{jA}); k_{iB}(\tau), k_{jB}(\tau\phi_{jB}))$: it depends on its own and its rival j 's marginal costs for *both* products. The key point is that i 's pass-through rate ρ_{iA} for product A will reflect any pricing spillovers emanating from product B —and this is captured by the formula in Proposition 2. On the cost side, multidimensional GLC allows for i 's unit cost to be identical across its offering, $k_{im}(\tau) = k_i(\tau)$, such that its cost shocks across components are also interlinked.

3.3. Special cases of multidimensional GLC. We next present examples of models that fit into the structure of multidimensional GLC. In each of these models, given A1M–A3M, GLC's core feature of supply linearity as per A4M is satisfied—and so Proposition 2 applies.

Multiproduct Cournot competition. Firm i offering consists of M products for which A1M–A3M are assumed to hold, and it chooses quantity for each product. For product $m \in M$, there is a set of N_m firms, including firm i itself. Firm i 's product $m \in M$ has a linear inverse demand curve given by: $p_{im} = \alpha_{im} - \sum_{j \in N_m} \beta_{jm} x_{jm} - \sum_{k \in M \setminus m} \sum_{j \in N_k} \gamma_{jk} x_{jk}$ which depends via $(\beta_{jm})_{j \in N_m}$ on firm i 's and rivals' supplies of product m as well as via $(\gamma_{jk})_{k \in M \setminus m, j \in N_k}$ on firm i 's and rivals' supplies of all other products apart from product m . GLC requires no assumptions about all of these parameters other than that they are constants; products can be arbitrarily substitutes or complements.¹⁹

Firm i 's overall profits $\Pi_i = \sum_{m \in M} (p_{im} - k_{im}) x_{im}$ and so the first-order condition for x_{im} is given by: $0 = \frac{\partial \Pi_i}{\partial x_{im}} = (p_{im} - k_{im}) - \left(\beta_{im} + \sum_{k \in M \setminus m} \gamma_{ik} \right) x_{im}$ where the latter term captures the cross-effect of the output quantity x_{im} of product m on the prices of other products. By inspection, this first-order condition for product $m \in M$ satisfies A4M with $\psi_{im} \equiv 1/(\beta_{im} + \sum_{k \in M \setminus m} \gamma_{ik})$. Therefore, this workhorse model is a member of multidimensional GLC and so Proposition 2 applies. As with baseline GLC, rivals' products could be entering and exiting—

¹⁹See Armstrong and Vickers (2018) and Nocke and Schutz (2018) for recent equilibrium analysis of multiproduct oligopoly pricing.

that is, the sets of firms $(N_m)_{m \in M}$ could vary with τ —as long as this does not affect $(\psi_{im})_{m \in M}$; insofar as this is relevant for firm i 's profit impact, it is captured by its product-level pass-through rates $(\rho_{im})_{m \in M}$.

Multiproduct Bertrand competition. The basic setup here is the same except that firm i now chooses prices. Firm i 's product m has a linear demand curve given by: $x_{im} = \alpha_{im} - \sum_{j \in N_m} \beta_{jm} p_{jm} - \sum_{k \in M \setminus m} \sum_{j \in N_k} \gamma_{jk} p_{jk}$ where the parameters $(\beta_{jm})_{j \in N_m}$ and $(\gamma_{jk})_{k \in M \setminus m, j \in N_k}$ are again taken to be constants.

Firm i 's overall profits $\Pi_i = \sum_{m \in M} (p_{im} - k_{im}) x_{im}$ so the first-order condition for p_{im} is: $0 = \frac{\partial \Pi_i}{\partial p_{im}} = x_{im} - (p_{im} - k_{im}) \beta_{im} - \sum_{k \in M \setminus m} (p_{ik} - k_{ik}) \gamma_{ik}$. Clearly, more work is needed to make this fit into the multidimensional GLC. A common simplifying assumption in the industrial-organization literature (see, e.g., Verboven, 2012) is that a firms' products have symmetric profit margins, with $(p_{im} - k_{im}) = \mu_i$ for all $m \in M$. In our setting, this also implies symmetric cost pass-through across products, with $\rho_{im} = \rho_i^M$ for all $m \in M$. While admittedly restrictive, this may still be a reasonable approximation of empirical behaviour in retail pricing where firms commonly set identical prices across different regional markets despite wide variation in demand conditions and competition (DellaVigna and Gentzkow, 2019). Recall that GLC is designed to admit such "behavioural" (and potentially irrational) pricing strategies. Given this, it follows that the first-order condition now satisfies A4M with $\psi_{im} \equiv (\beta_{im} + \sum_{k \in M \setminus m} \gamma_{ik})$. Therefore, a partially-symmetric linear version of multiproduct Bertrand competition is also a member of multidimensional GLC.

Multiproduct quality competition: Upgrades approach. We next consider a simplified version of the upgrades approach (Johnson and Myatt, 2003, 2006; Johnson and Rhodes, 2018) to multiproduct quality competition. Each of $n \geq 2$ firms offers two product qualities: low-quality q_1 and high-quality q_2 , where $q_2 > q_1$. We consider a special case of the model that yields linear demand structures—so that GLC's A4M holds.

A consumer of type θ has multiplicative willingness-to-pay $v(\theta, q) = \theta q$ 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 Y_1^i denote the combined number of low- and high-quality units produced by firm i , let Y_2^i be the number of high-quality units, so that $(Y_1^i - Y_2^i)$ is the number of low-quality

units. Similarly, let $Y_1 = \sum_{i=1}^n Y_1^i$ be the industry supply of both qualities and $Y_2 = \sum_{i=1}^n Y_2^i$ that of high-quality. In the upgrades approach, the former is interpreted as the number of “baseline” units while the latter is the number of “upgrades” (from low to high quality). Due to multiplicative preferences and uniform types, both demand curves $p_B(Y_1) = (1 - Y_1)q_1$ and $p_U(Y_2) = (1 - Y_2)(q_2 - q_1)$ are linear. The total price of the high-quality product is the baseline price plus the upgrade price, $p_B + p_U$.

We assume that multidimensional GLC’s A1M–A3M are satisfied for both firm i ’s baseline and upgrade products. Firm i ’s unit cost for the low-quality product is $k_B^i(\tau) = c_B^i(\tau) + \tau z_B^i(\tau)$ while it is $k_U^i(\tau) = c_U^i(\tau) + \tau z_U^i(\tau)$ for the upgrade, where the corresponding regulatory intensities, $z_B^i \equiv e_1^i/Y_1^i$ and $z_U^i \equiv e_2^i/Y_2^i$, are both optimally chosen given τ . Hence, by the envelope theorem, $dk_m^i(\tau)/d\tau = z_m^i(\tau)$ for $m \in \{B, U\}$.

Firm i ’s overall profits are therefore given by:

$$\begin{aligned} \Pi^i(\tau) &= [p_B(Y_1) - k_B^i(\tau)](Y_1^i - Y_2^i) + [p_B(Y_1) + p_U(Y_2) - k_B^i(\tau) - k_U^i(\tau)]Y_2^i \\ &= \underbrace{[p_B(Y_1) - k_B^i(\tau)]Y_1^i}_{=\Pi_B^i} + \underbrace{[p_U(Y_2) - k_U^i(\tau)]Y_2^i}_{=\Pi_U^i}, \end{aligned}$$

where it separately chooses Y_1^i and Y_2^i (subject to $Y_1^i \geq Y_2^i$). It is easy to check that the two first-order conditions for Y_1^i and Y_2^i satisfy multidimensional GLC’s A4M, with $\psi_B^i = 1/q_1$ and $\psi_U^i = 1/(q_2 - q_1)$. The upgrades approach has the general feature that the baseline units and upgrades are neither substitutes nor complements. A substantive restriction is that regulation τ does not alter the quality levels q_1, q_2 offered by firms. Hence Proposition 2 pins down the profit impact, with pass-through defined as $\rho_m^i \equiv [dp_m(\tau)/d\tau]/[dk_m^i(\tau)/d\tau]$ for product $m \in \{B, U\}$.

As in the previous extensions, it is worth stressing that this follows directly from firm i ’s first-order conditions. There are no restrictions on its rivals’ technologies or behaviour or on the equilibrium concept—though the upgrades approach does involve assumptions about the maximizing behaviour of consumers.

Multimarket network competition. Firms often compete by selling their product across multiple markets. Recent work in the networks literature has sought to understand how strategic interaction and network structure affect market outcomes (Bramoullé et al., 2014). We here show how a simple linear model of network competition, used by Elliott and Galeotti (2019) in

the context of antitrust investigations, fits into the GLC framework.

The model has $n \geq 2$ firms selling a product into a set of $M \geq 2$ different consumer markets. Firm i has production capacity K_i and sells $x_{im} = \sigma_{im}K_i$ units into market l , where $\sum_{m \in M} \sigma_{im} = 1$, so $\sigma_{im} \in (0, 1)$ reflects i 's presence in market m . Given firms' overall capacities \mathbf{K} , the price in market m is determined by a linear demand curve $p_m(\mathbf{K}) = \alpha_m - \beta X_m$, where $X_m = \sum_{i=1}^n \sigma_{im}K_i$ and the demand parameter α_m may vary across markets. Firm i 's overall sales to all M markets are $x_i = \sum_{m \in M} x_{im} = K_i$, that is, it sells up to capacity. On the production side, the good sold by firm i has the same cost structure across consumer markets. We assume that its unit cost satisfies A1M–A3M which, given symmetry, becomes isomorphic to baseline GLC's A1–A3—with an optimized unit cost $k_i(\tau)$ and regulation intensity $z_i(\tau) \equiv e_i(\tau)/x_i(\tau)$ for each market $m \in M$

Firm i chooses its capacity K_i to maximize its overall profits $\Pi_i(\mathbf{K}; \tau) = \sum_{m \in M} [p_m(\mathbf{K}) - k_i(\tau)]\sigma_{im}K_i$, which yields the first-order condition: $\sum_{m \in M} [p_l(\mathbf{K}) - k_i(\tau)]\sigma_{im} - \beta \sum_{m \in M} \sigma_{im}^2 K_i = 0$ that defines i 's optimal capacity as a function of the network structure as well as demand and cost conditions.²⁰ Let $\bar{p}_i \equiv \sum_{m \in M} \sigma_{im}p_m$ be i 's weighted-average price across its markets, and use this to rewrite i 's first-order condition as: $\bar{p}_i(\tau) - k_i(\tau) = (\beta \sum_{m \in M} \sigma_{im}^2) K_i(\tau)$, now also making explicit dependencies on regulation τ . This supply schedule has exactly the same form as baseline GLC's A4, with a constant slope parameter, $\psi_i = 1/(\beta \sum_{m \in M} \sigma_{im}^2)$, that reflects the production network and demand conditions. Proposition 2 follows by letting $\rho_i^M \equiv [d\bar{p}_i(\tau)/d\tau]/[dk_i(\tau)/d\tau]$ denote i 's multidimensional rate of cost pass-through (and noting that the weights $\omega_{im} = \sigma_{im}$ for all $m \in M$).

This example shows how our GLC result on firm-level pass-through applies to a multimarket setting. A conceptual difference compared to the networks literature is that Proposition 2 is derived by making assumptions only about firm i itself. There are no assumptions as such on the technology or behaviour of i 's rivals; all salient characteristics of the network are captured in i 's pass-through rate ρ_i^M . A substantive restriction, also employed in the linear-quadratic games that are widely studied in the networks literature, is that the network structure \mathbf{g} is

²⁰While this network model clearly has a multidimensional flavour, the combination of symmetric costs and a single choice variable mean that its structure effectively boils down to baseline GLC.

fixed; in our case, it does not vary with regulation τ .²¹

Oligopolistic price discrimination. In practice, consumers often pay different prices for the same good; for example, airlines sell tickets at different prices depending on when a customer buys. We next show how Proposition 2 applies to a model of oligopolistic price discrimination (Hazledine, 2006; Kutlu, 2017).

Consumers have unit demand for a homogeneous product (e.g., an economy flight from A to B), with a distribution of $v(X) = 1 - X$ (so the X^{th} keenest consumer has value $1 - X$). There are $M \geq 2$ price buckets with class m priced at $p_m = 1 - X_1 - \dots - X_m$, where X_m is the number of units sold in class m , so earlier buckets have higher prices. This demand system is an oligopolistic extension of a setup in Varian (1985); the number of buckets measures the extent of price discrimination. Each firm $i \in N$ chooses how much of each bucket to supply $(x_{im})_{m=1}^M$.

On the production side, the good sold by firm i has the same cost structure across price buckets. Its unit cost satisfies A1M–A3M which, given symmetry, becomes isomorphic to baseline GLC’s A1–A3—with an optimized unit cost $k_i(\tau)$ and regulation intensity $z_i(\tau) \equiv e_i(\tau)/x_i(\tau)$ for each bucket $m \in M$. For expositional simplicity, we here also assume that all firms have identical unit costs.²² Firm symmetry implies that $x_{im}/x_i = X_m/X$ (for all $i \in N$). Given regulation τ , firm i ’s overall profits are given by $\Pi_i = \sum_{m=1}^M \Pi_{im} = \sum_{m=1}^M [p_m - k_i(\tau)]x_{im}$

This model of third-degree price discrimination has a complex demand structure: the price for one bucket depends on firms’ output choices for all other buckets. However, at the symmetric Nash equilibrium, any firm i ’s optimality condition for price bucket m can be written as $[p_m - k_i(\tau)] = h_m(m, M, n)x_{im}$. The key feature is that the proportionality term $h_m(m, M, n)$ depends only on market structure and the number of price buckets. Therefore, multidimensional GLC’s A4M is satisfied and Proposition 2 applies, as long as regulation τ does not alter M or n . The weights underlying multidimensional pass-through, $\rho_i^M(\tau; \Phi) \equiv \sum_{m \in M} \omega_{im}(\tau)\rho_{im}(\tau; \Phi)$, here satisfy $\omega_{im}(\tau) \equiv e_{im}(\tau)/e_i(\tau) = X_m(\tau)/X(\tau)$ given that regulatory intensity is the same across

²¹The networks literature characterizes Nash equilibrium based on the first-order conditions of all firms. It would typically define $g_{ij} \equiv (\sum_{l=1}^m \sigma_{il}\sigma_{jl})/\sum_{l=1}^m \sigma_{il}^2$ as a key measure of the closeness of competition between firms i and j across the m markets so that the $n \times n$ matrix \mathbf{g} represents the network of interaction. (If i and j serve distinct subsets of markets, then $\sigma_{il}\sigma_{jl} = 0$ for all l and so also $g_{ij} = 0$; if they serve identical markets, then $\sigma_{il} = \sigma_{jl}$ for all l and so $g_{ij} = 1$.) The central results in the literature use different network measures—notably Bonacich centrality—to characterize Nash equilibrium (Ballester et al., 2006; Bramoullé et al., 2014).

²²The main points about GLC also obtain in richer setups with asymmetric costs among firms and conduct other than Cournot-Nash.

price buckets and firms are symmetric.

The central result in this strand of literature is that price discrimination does not change the average price paid by consumers. Greater price discrimination (higher M) raises (lowers) the prices paid by high-value (low-value) consumers but, due to the linear model structure, leaves the average price unchanged.²³ In particular, the average price $p_{\text{ave}}(M) \equiv \sum_{m=1}^M \frac{X_m}{X} p_m = p_{\text{ave}}(1)$ for all $M \geq 2$, where $p_{\text{ave}}(1)$ is the uniform Cournot price.

This result also has significant implications for pass-through. It implies immediately that the change in price due to regulation satisfies $\frac{d}{d\tau} p_{\text{ave}}(M) = \frac{d}{d\tau} \sum_{m=1}^M \frac{X_m(\tau)}{X(\tau)} p_m(\tau) = \frac{d}{d\tau} p_{\text{ave}}(1)$ for all $M \geq 2$. Firm i 's rate of cost pass-through for price bucket m is defined as $\rho_{im} \equiv [dp_m(\tau)/d\tau]/[dk_i(\tau)/d\tau]$ and multidimensional pass-through $\rho_i^M(\tau; \Phi) \equiv \sum_{m \in M} \omega_{im}(\tau) \rho_{im}(\tau; \Phi)$. The weights satisfy $\omega_{im}(\tau) = X_m(\tau)/X(\tau)$ so $\rho_i^M(\tau; \Phi) = [\frac{d}{d\tau} p_{\text{ave}}(M)]/[dk_i(\tau)/d\tau]$ is also constant with respect to the extent of price discrimination M . Greater price discrimination affects relative prices and relative pass-through—but the average rate of pass-through is unaffected.

In sum, this extension (i) shows how GLC's main results extend to a setting with price discrimination, and (ii) provides a microfoundation, within the context of GLC, for working with cost pass-through based on an average of dispersed market prices.

4. EMPIRICAL ANALYSIS OF COST PASS-THROUGH FOR US AIRLINES

In the rest of the paper, we demonstrate the applicability and usefulness of the GLC framework by estimating the impact of introducing carbon pricing to the domestic US airline market. The empirical implementation of GLC is based on three main steps. First, the scale τ and scope Φ of the cost shift needs to be identified and a corresponding estimate of firm i 's rate of cost pass-through obtained. Second, GLC's assumptions must be deemed a reasonable approximation to firm i 's production technology and competitive environment. Third, the appropriate GLC results must be used to calculate the impacts of regulation. Given prior knowledge of firm-level pass-through, profit impacts can be estimated via GLC using only cross-sectional data for the firms of interest. If, as we do here, firm-level pass-through needs to be estimated, then the data requirement increases to include data on prices and quantities of the input factor, plus an instrument for the cost shift.

²³Price discrimination still benefits firms because sales expand to otherwise excluded low-value consumers.

4.1. Background on aviation and climate change policy. Airlines currently produce around 2.5% of global CO₂ emissions (McCollum et al., 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. 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.²⁴

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 \$30/tCO₂, aviation’s 2013 emissions of 120 million tCO₂ would have had a value of \$3.6 billion. We study the domestic market because international aviation is regulated under a separate organization and set of agreements.²⁵

There is a rich literature that explores many important features of US aviation markets, principally the presence of market power and firm heterogeneity. The number of firms competing on a given route is typically small, and it is widely acknowledged that despite deregulation airlines continue to exercise market power (Brander and Zhang, 1990; Borenstein and Rose, 2007; Ciliberto and Tamer, 2009). Over the period we study, US aviation was dominated by four large legacy carriers (American, Delta, United and US Airways) and a large low-cost carrier, Southwest Airlines. Legacy airlines were established on interstate routes 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). A recurring theme in the literature is the importance of this heterogeneity and the central role played by Southwest, sometimes labelled the “Southwest is special” effect. From the 1980s onwards, Southwest’s low costs, innovative business model and generally astute management

²⁴Aviation is already subject to carbon pricing in some international jurisdictions such as the EU’s emissions trading scheme.

²⁵The 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.

have made it an especially disruptive and profitable competitor, with large impacts on incumbent prices and profits (Borenstein, 2011; Morrison, 2001; Borenstein and Rose, 2007; Goolsbee and Syverson, 2008; Ciliberto and Tamer, 2009).²⁶ We follow this literature by ensuring potential heterogeneity between Southwest and the legacy carriers is captured in our analysis.

4.2. Description of the data. As is standard in the airline literature, we take a product to be a seat on a flight by airline i on route j at time t . Our dataset is a panel of price and cost data for airlines over the period 2004Q1–2013Q4.²⁷ 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. All monetary quantities are in real 2013Q4 USD. We restrict our attention to the four large legacy carriers (American, Delta, United and US Airways) and one large low-cost carrier (Southwest) operating throughout the period.²⁸ The resulting dataset is an unbalanced panel, with $N = 1334$ carrier-routes and $T = 40$ quarters, a total of 36,650 observations. The routes in our sample make up 27% by revenue of all domestic US aviation activity over the period.

We construct the 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.²⁹ 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. Following much of the airlines literature, we exclude indirect flights.³⁰ Finally, we exclude carrier-routes that had fewer than 1,000 passengers per quarter (i.e. 83 per week). For the resulting observations we calculate p_{ijt} , an average of all fares

²⁶The “Southwest is special” effect is widely recognised beyond the economics literature. See, for example, Heskett and Sasser (2010) and Tully (2015) for business discussions.

²⁷The start year was chosen to exclude the effects of 9/11 from our dataset. The end date avoids the 2015 “mega merger” between American Airlines and US Airways.

²⁸There are now many more low-cost carriers, but these were either small or non-existent at the start of our period, so we do not include them.

²⁹We 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. Our multidimensional GLC results, applied to the upgrades approach (Johnson and Myatt, 2003, 2006), 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.

³⁰Indirect 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 and Rose, 1994; Goolsbee and Syverson, 2008; Gerardi and Shapiro, 2009).

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 .³¹ This is the total fuel cost for a given flight divided by the number of passengers on that flight (Appendix B describes this procedure in more detail). Given there is currently no variation in carbon prices for US airlines, our empirical approach is to estimate fuel cost pass-through – the dollar increase in prices following a dollar increase in fuel costs. There is a constant relationship whereby 1 gallon of jet fuel produces 0.00957 tons of CO₂ when burned; for an airline, fuel and carbon costs are therefore equivalent and pass-through rates should be the same.³²

The value of fuel cost 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, 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 instruments 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.³³

Table 2 presents descriptive statistics on airlines’ prices, costs and other variables related to competition and environmental performance. The four legacy carriers are grouped together (Table C1 in Appendix C gives descriptive statistics by carrier and confirms the similarity of the legacy carriers compared to Southwest). Southwest tends to fly larger numbers of passengers on shorter routes than the legacy carriers; it charges lower prices and has lower fuel costs and

³¹The overlap between the DB1A and T-100 is good but not perfect (see Goolsbee and Syverson, 2008, for a fuller discussion). Merging with data from T-100 results in around 10% of DB1A revenue being dropped.

³²A similar approach of using variation in other input costs to estimate the impact of future environmental costs is also taken by Miller et al. (2017). In related work, Ganapati et al. (2016) estimate the pass-through of energy input prices across six US manufacturing industries while Bushnell and 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.

³³We do not have detailed data on the precise extent of hedging by each carrier at each point in time. Turner and Lim (2015) document and analyse the different hedging strategies of US airlines, and the different effective fuel prices that result.

emissions. Revenue and numbers of competitors are broadly similar across routes.

Figure 2 shows trends over the period for each carrier type. Figure 2(a) compares Southwest’s average per-passenger fuel cost with the spot price of jet fuel. They track each other reasonably closely, with a lag indicating the presence of hedging, which also smooths out the peak and trough from the 2008 price spike. Note also the substantial variation in fuel costs over the period. Figure 2(b) plots average ticket prices (left axis) against per-passenger fuel costs (right axis) for Southwest. As expected, there is a positive correlation. Figure 2(c) shows per-passenger fuel costs for the legacy carriers and how they compare to the spot price of jet fuel. Fuel costs follow spot prices more closely than for Southwest, consistent with a more limited use of hedging. Figure 2(d) shows ticket prices and fuel costs for the legacy carriers; as for Southwest, these appear to be closely related.

4.3. Step 1: Estimating pass-through. The first step in simulating the effect of carbon pricing in the US aviation industry is to identify the scale τ and scope Φ of regulation and obtain an estimate of firm i ’s rate of cost pass-through. We assume a federal carbon price that is equally applied across all airlines ($\Phi = 1$). To estimate firm i ’s rate of cost pass-through, we follow the recent pass-through literature (e.g., Fabra and Reguant, 2014; Atkin and Donaldson, 2015; Stolper, 2016; Miller et al., 2017) and use as our baseline specification a panel data regression of quarterly average prices p_{ijt} on average per-passenger fuel costs k_{ijt} . We allow for pass-through heterogeneity by interacting costs with carrier identity – either Southwest or a legacy carrier – and estimate the following equation:

$$p_{ijt} = \rho_S k_{ijt} \cdot S_i + \rho_L k_{ijt} \cdot L_i + X'_{ijt} \beta + \lambda_t + \eta_{ij} + \epsilon_{ijt}. \quad (7)$$

The parameter of interest is ρ , the average rate of cost pass-through across routes. By interacting fuel cost k_{ijt} with dummies S_i and L_i that are equal to 1 for Southwest and legacy carriers, respectively, we obtain pass-through rates ρ_S and ρ_L for the two carrier types. Controls (X_{ijt}), time effects (λ_t), fixed effects (η_{ij}), and residuals (ϵ_{ijt}) are explained below.

The per-passenger fuel cost k_{ijt} in equation (7) is potentially endogenous. It depends in its denominator on the number of passengers flying in quarter t , which in turn will generally be an outcome of the price p_{ijt} . To address this, we use the spot price of jet fuel as an instrument for

fuel cost k_{ijt} .³⁴ 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 restriction. To accommodate potential hedging, we include three lags of spot prices. We also include a term interacting distance with fuel costs, in order to capture the differential impact of spot fuel prices on the dollar fuel cost of flights of different lengths. The first stage regressions are given by:

$$k_{ijt} \cdot a_i = \sum_{q=0}^3 \gamma_q f_{t-q} \cdot a_i + \sum_{q=0}^3 \delta_q f_{t-q} \cdot a_i \cdot d_{ij} + X'_{ijt} \beta + \mu_t + \theta_{ij} + \omega_{ijt} \quad \text{for each } a_i \in \{S_i, L_i\} \quad (8)$$

where f_t is the spot fuel price, d_{ij} is the route distance, and the remaining controls are the same as in Equation (7).

Equations (7) and (8) include a vector of controls X_{ijt} which capture changes in supply and demand: (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 use the average GDP growth of the two states at either end of the route. (ii) We construct an index of the main non-fuel costs, at the carrier level, principally made up of labour and aircraft maintenance costs. (iii) We include the number of all competing legacy carriers on route j . (iv) We also include the number of all competing low-cost carriers, as these may have a different impact to other competitors. (v) We include carrier size, the total number of passengers travelling on all of airline i 's routes in a given quarter, as a measure of overall demand for airline i 's product. Time effects are controlled for using indicators for each year-quarter t and η_{ij} and θ_{ij} capture carrier-route level fixed effects. Standard errors are clustered at the carrier-route level. Finally, we weight observations by carrier-route emissions, so smaller routes do not disproportionately affect the resulting carrier-level pass-through estimates and corresponding to our definition of multidimensional pass-through.³⁵

³⁴We use jet fuel price data from Bloomberg (JETINYPR index, New York Harbor 54-Grade Jet Fuel).

³⁵Our specification does not directly include competitors' costs, as this makes estimation infeasible because of a curse of dimensionality problem. The literature on pass-through estimation takes a variety of approaches to deal with this problem, for example by imposing functional form or other restrictions to reduce the number of parameters to estimate (see, for example, Miller et al., 2017). 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's costs. The substantive assumption is that the relationship between own and rivals' costs is unchanging over the period we study. If this holds, then our pass-through rates capture both own-cost and rival-cost effects. In this respect our approach is similar to both Atkin and Donaldson (2015) and Stolper (2016).

4.4. Estimation results. Table 3 shows our main estimation results. Column 1 presents the OLS coefficients. Interestingly, the pass-through is not only significantly higher for Southwest than the legacy carriers, but also these are on either side of 1. Column 2 provides the baseline results correcting for the cost endogeneity through 2SLS. Southwest’s pass-through rate is 1.42 and significantly greater than 1 (p-value: 0.001), whereas the legacy carriers’ pass-through is 0.66 and significantly smaller than 1 (p-value: 0.001). This means that when Southwest’s costs rise, their rise in equilibrium price more than offsets this. The opposite is true for the legacy carriers. Both coefficients are higher than their counterparts in column 1, confirming the OLS downward bias, while at the same time revealing the same qualitative pattern. To the best of our knowledge, the finding of pass-through heterogeneity between Southwest and legacy carriers in the US airline industry is novel. We document, therefore, a new dimension in which the “Southwest is special” result operates.

The rest of the control coefficients generally have the expected signs. We see that greater competition lowers prices, and that competition from low-cost carriers has a large additional effect. Non-fuel costs raise prices as expected. GDP growth and carrier size seem to have a negative relationship with price. The first stage results are also as expected: both the spot price of jet fuel and its interaction with distance have a positive relationship with per-passenger fuel costs, and the Sanderson-Windmeijer F-statistics provide very strong evidence that the instruments are relevant in all cases, in line with past pass-through literature.

Column 3 shows the results of estimating the same equations on the subsample of the full dataset that gives a balanced panel. These results are for the ‘stable’ routes only—those that were operated continuously by the same airlines throughout the period. The near-identical results for the balanced and unbalanced panel suggest that both newly opened routes and routes shortly to be discontinued appear to have the same pass-through as the stable routes. This is encouraging for the external validity of our estimates.

4.5. Pass-through heterogeneity. What explains the heterogeneity in pass-through rates? To answer this we can decompose the difference in pass-through rates into three factors: route portfolio, production costs, and demand asymmetries. First, from the descriptive statistics in Table 2 it is clear that Southwest operates considerably shorter routes than the legacy carriers.

Short-haul flights are likely to have systematically different characteristics: on the demand side, there are more close substitutes (such as car, bus or rail travel); on the supply side, entry may be more or less difficult, so there may be differences in market power (e.g., Brander and Zhang, 1990; Berry and Jia, 2010). To quantify the importance of route portfolios, we present in column 4 of Table 3 the results using only routes common to both Southwest and the legacy carriers.³⁶ We find that Southwest’s pass-through rate falls significantly, to 0.98, while the legacy carriers’ pass-through remains statistically unchanged, at 0.52. Hence, 39% of the difference in pass-through rates between Southwest and legacy carriers (Table 3, column 2) can be attributed to route portfolio differences.³⁷

Second, we quantify the importance of production cost heterogeneity. Southwest has lower average fuel costs on common routes: \$30.37 per passenger compared to \$40.59 for the legacy carriers. This is principally due to use of newer aircraft and more efficient seating configurations. Absent any demand asymmetries, prices would have been equal and from equation (4) the pass-through rates of any two firms (i and j) are related only by their relative cost shocks, $\bar{p}_j/\bar{p}_i = \Delta k_i/\Delta k_j$. Using the fuel cost figures for the common routes, Southwest’s superior fuel efficiency can explain 23% of the original difference in baseline pass-through rates.³⁸

Third, the remaining 38% of the pass-through differential can be attributed to demand-side asymmetries between carriers, based on their differentiated-product offering. An alternative, but perhaps less likely explanation, is that these demand asymmetries are driven in addition by differences in competitive conduct across common routes.³⁹

4.6. Further results and robustness. We have explored a large number of robustness tests and alternative specifications, and find that our central result of pass through greater than 1 for Southwest and less than 1 for the legacy carriers is robust. We summarize the results here, while Appendix B contains further details. First, in Appendix B.2 we re-estimate our baseline results but allowing for heterogeneity at the individual airline level, rather than grouping the

³⁶To achieve a fully like-for-like comparison of pass-through rates we run an unweighted regression.

³⁷The difference in coefficient across all routes is $1.418 - 0.661 = 0.757$, whereas the difference across common routes is $0.983 - 0.518 = 0.465$, hence $(0.757 - 0.465)/0.757 = 0.39$.

³⁸All else equal, we would expect Southwest’s predicted pass-through rate to be 0.692, which explains $(0.692 - 0.518)/0.757 = 0.23$ of the difference in coefficients in Table 3 (column 2).

³⁹To add together the above elements of this back-of-the-envelope decomposition, we implicitly require the ratio of fuel efficiencies $\Delta k_i/\Delta k_j$ between Southwest and the legacy carriers to be the same for the whole sample as it was in the common routes.

legacy carriers together. Table B2, column 2, reveals that Southwest remains the only carrier to have a pass-through rate significantly greater than 1. Second, in Appendix B.3 we explore a larger set of control variables that address additional demand- and supply-side factors that could be relevant to airlines' pricing choices (GDP per capita, population, network density, potential entry by Southwest à la Goolsbee and Syverson (2008) and the effect of bankruptcy). The pass-through rates in Table B4 remain qualitatively and statistically unchanged. Third, in Appendix B.4 we present results of estimating Equation (7) on a route-by-route level rather than as a panel, in a way closely related to the Mean Group estimator in Pesaran and Smith (1995) and the estimation strategy used in Atkin and Donaldson (2015). The principal advantage of this approach relative to panel regressions is that it allows for greater heterogeneity across routes within an airline, at the cost of less comprehensive time controls. The results in Table B5 are again similar to our baseline.

5. CALCULATING THE IMPACTS OF CARBON PRICING USING GLC

In this section, we combine our pass-through estimates with GLC theory to calculate the predicted impact of a carbon price on US airlines.

5.1. Step 2: Verifying GLC's assumptions. The second step, before applying GLC's results, is to verify that the underlying assumptions are a reasonable approximation to firm i 's production technology and competitive environment. GLC requires no specific assumptions about demand functions, the mode of competition, competitors' technologies and strategies, rationality or equilibrium. Many of these factors are unknown for airlines, and there will be considerable intra-industry heterogeneity. As each airline operates in multiple markets, we use multidimensional GLC and next discuss its assumptions.

Input price-taking (A1M) is a reasonable assumption in that any 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.

Cost-minimizing inputs (A2M) also appears reasonable for airlines. Fuel costs are often an airline's largest cost, amounting to 20–50% of its total cost base (Zhang and 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.⁴⁰

Constant returns to scale (A3M) is a more substantive assumption, though it is standard in much of the airlines literature.⁴¹ The evidence on whether it holds empirically is inconclusive: while some studies estimate modest scale economies others find no such evidence (Zhang and 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.⁴² Note also that the presence of fixed costs is not an issue for the application of GLC.

Linear product market behaviour (A4M) is a key assumption underlying GLC. A direct test of whether A4M holds in a given setting would require full information on marginal costs. Obtaining such information is known to be extremely difficult in most markets. The alternative of estimating marginal costs through a structural model (for example, Bresnahan, 1989; Berry, et al., 1995; Nevo, 2001) would require a whole set of additional assumptions that defeats the purpose of using the GLC framework in the first place. Instead, an indirect way to test this assumption is by providing evidence that demand does not exhibit significant non-linearities; as explained above, given A1M-A3M, standard models of imperfect competition then imply that A4M holds. We conduct two relevant empirical tests. First, we estimate pass-through on monopoly routes using our baseline model (7). The estimated pass-through is 0.6 (Table B6 in Appendix B) and statistically not different from 0.5, which is the pass-through rate for a monopolist facing linear demand (Bulow and Pfleiderer, 1983). This is in line with Genakos and Pagliero (2019), who find a similar result for monopolists in the retail petroleum market. This result provides direct evidence that demand convexity is not significant for monopoly routes.

⁴⁰Other 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 pass-through results may give less good predictions of the impact of future carbon pricing.

⁴¹Brander and Zhang (1990) discuss how to conceptualise constant marginal costs in the case of airlines; Berry and Jia (2010) estimate marginal costs which are constant for a given vector of route characteristics.

⁴²There is stronger evidence for economies of scope: a higher network density of its route portfolio can confer a competitive advantage on an airline—but this is not inconsistent with GLC theory. As is common in the airlines literature, our empirical analysis does not account for potentially complex network effects with other routes (see, e.g., Ciliberto and Tamer, 2009).

Second, given that GLC allows for non-linearities in rivals' prices (see Remark 4), we directly test for the existence of non-linear demand terms in own price in duopoly and triopoly markets. In both cases, we fail to reject the null of insignificant coefficients (details are in Appendix B.5). In sum, A1M-A3M combined with evidence of no significant demand non-linearities provide empirical support to maintain A4M in our application.

5.2. Step 3: Profit impacts of an exogenous carbon price. Given that A1M-A4M hold, we can apply multidimensional GLC's Proposition 2 to estimate the profit impacts of a carbon price. We begin by examining the impact of an exogenous carbon price of \$30/tCO₂, which is roughly the social cost of carbon (SCC) reported by Nordhaus (2017) and calculated by the US government.⁴³ Importantly for the external validity of our results, the wide variation of jet fuel prices over our period exceeds by a considerable margin the variation in costs that a \$30/tCO₂ carbon price would induce. Hence, our simulated carbon-price shock lies within the range of fuel cost shocks that airlines responded to over the sample period.

Using our baseline pass-through estimates (Table 3, column 2) and Proposition 2(c), the impact of a \$30/tCO₂ carbon price is summarised in Table 4, columns 1-3.⁴⁴ The legacy carriers' profits are expected to fall by approximately \$234 million, whereas Southwest's profits increase by \$98 million. As a percentage of revenue, these amounts correspond to a 1.60% decrease for the legacy carriers and a 1.55% increase for Southwest. The intuition is that, on average, legacy carriers' pass-through is less than 100%, so their profit margins decline as costs increase; by multidimensional GLC's A4M, this leads to a fall in their sales and profits. For Southwest, this logic applies in reverse and so its profits increase due to carbon regulation. The overall effect on the airline industry is a \$136 million (or alternatively a modest 0.65%) decline in profits. This is much less than the \$463 million profit decline from a naive "static" calculation

⁴³The US Interagency Working Group on Social Cost of Greenhouse Gases (IAWG 2016) has compared and averaged the results of the major models that are used to estimate the SCC. The results for the end of our period of study are around \$30/tCO₂.

⁴⁴Proposition 2(c) also requires two auxiliary conditions: (i) the cost shock due to regulation is "small" relative to product price and (ii) there is limited scope for factor substitution. First, from Table 2, emissions costs at \$30/tCO₂ are about 3% of the average ticket price for both Southwest and legacy carriers. Second, clearly there is limited scope for large-scale jet fuel substitution. However, in practice, airlines can increase fuel efficiency by adjusting various aspects of a flight (e.g., route and altitude to minimize the impact of wind, holding patterns and taxiing between runway and terminal, etc.). We estimate empirically that a 1% increase in carbon prices decreases the emissions' intensity by -0.03% for legacy carriers to -0.01% for Southwest (results not reported here, available on request). In sum, we interpret both auxiliary conditions of Proposition 2(c) as being met (more details in Appendix B.6).

($= -\sum_{i=1}^n (\bar{\tau} - \underline{\tau}) e_i(\underline{\tau})$). GLC’s results identify the winners and losers from carbon pricing, and therefore also inform us about the likely political game to be played.

While our finding of pass-through heterogeneity in the airline industry is novel, the conclusions on the calculated profit impact via GLC are consistent with several findings in the literature. Goolsbee and Syverson (2008), Ciliberto and Tamer (2009) and Berry and Jia (2010) variously point to Southwest being more efficient, better able to cope with shocks, or especially threatening to its rivals (the “Southwest is special” effect discussed above). Gaudenzi and Bucciol (2016) report jet fuel price rises are associated with significantly more negative stock-market returns for legacy carriers than for Southwest. While these authors stress differences in hedging strategies, our findings offer a new explanation for their results.

5.3. Consumer surplus and social welfare. We now extend our analysis by showing how the multidimensional GLC structure—combined with some additional assumptions—yields results on social welfare (see Appendix C.1 for details).

To be able to make statements about welfare, we introduce two additional assumptions. A5M assumes that aggregate consumer surplus S can be written in terms of a gross consumer utility function $V(\cdot)$ that depends on the quantities of the products offered by all firms in the market. Our formulation continues to allow for a wide range of demand-side properties such as substitutes and complements, differentiated products, and so on; we do not require specific functional-form assumptions on $V(\cdot)$. A6M assumes that environmental damages $D(E)$ depend on industry-wide emissions E . In other words, which firms produce emissions does not matter; only the aggregate is relevant. This is the appropriate setup for climate change.

Social welfare can therefore be written as $W(\tau) = S(\tau) + \Pi(\tau) + \tau E(\tau) - D(E(\tau))$, where the third term is government revenue.⁴⁵ To obtain welfare results, which depend on industry wide-profits Π , we now assume that multidimensional GLC’s A1M–A4M hold for each firm in the industry.

We thus obtain two further sufficient-statistics results summarized in Proposition 3. First, the change in consumer surplus (ΔS) is driven by the rate of industry-average multidimensional pass-through $\tilde{\rho}^M \equiv \sum_{i=1}^n \frac{e_i}{E} \rho_i^M$. Second, the change in social welfare (ΔW) depends on $\tilde{\rho}^M$ as

⁴⁵Government revenue from carbon pricing is a transfer from producers.

well as the induced change in industry-wide emissions ($dE(\tau)/d\tau$). This shows how pass-through is also critical for welfare analysis under GLC.

Column (3) of Table 4 applies these results to US airlines. Based on our carrier-level pass-through estimates, we obtain an industry-average pass-through rate $\tilde{\rho}^M = 0.853$. This yields that consumer surplus declines by \$395 million at our \$30/tCO₂ carbon price. Conversely, we find that social welfare rises by \$49 million: the reduction in consumer surplus and industry profits is outweighed by environmental gains.⁴⁶ However, our confidence levels do not allow us to rule out the possibility that total welfare declines at the lower end of our pass-through estimates (for which the decline in industry profits is larger).

5.4. Political economy and an endogenous carbon price. So far we have focused on a regulation τ that is exogenous. However, one could argue that regulation cannot be treated as a “manna from heaven”, especially for an influential industry like US airlines. We therefore show in Proposition 4 (see Appendix C.2 for details), for the case of an emissions tax, how regulation can be endogenized using multidimensional GLC’s A1M–A6M. This analysis brings together two strands of prior research: (1) an influential literature following Grossman and Helpman (1994) in which firms lobby a government “for sale”; (2) a classic literature following Buchanan (1969) on the optimal design of emissions taxes under imperfect competition in the product market.⁴⁷ In the model, the government cares about social welfare—which includes the social cost of carbon emissions—but also gains utility from “contributions” by corporate lobbyists. Proposition 4 derives the optimal carbon price $\tau^*(\lambda)$ as a function of the Grossman-Helpman lobbying parameter λ . It generalizes existing literature to richer modes of imperfect competition and provides a unifying result in terms of the industry-average of firm-level pass-through rates.

We apply this result to US airlines. Both distortions push downwards the equilibrium carbon price: (1) as industry-average pass-through $\tilde{\rho}^M = 0.853$ lies below 100%, the industry collectively opposes a carbon tax and lobbies the government to water it down;⁴⁸ (2) to mitigate

⁴⁶To obtain this result, we estimate the induced change in industry-wide emissions using our airlines data, confirming that carbon pricing reduces aggregate emissions, $dE(\tau)/d\tau < 0$. A detailed description is given in Appendix B.7.

⁴⁷By assumption, the government does not have access to another policy instrument (such as a price control) to directly address market power.

⁴⁸The exception is that, by GLC’s logic, Southwest should individually endorse the carbon tax.

deadweight losses from market power, the government lessens the increase in product prices by softening the carbon price. Taking the social marginal damage as \$30/tCO₂, and calibrating λ using prior literature, we find an endogenous carbon price of \$17.68/tCO₂—around 41% lower than the Pigouvian benchmark. The profit impacts of regulation, shown in columns (4)-(6) of Table 4, are therefore less pronounced for the endogenous carbon price: a loss of only 0.94 percentage points for the legacy carriers and a gain of only 0.91 percentage points for Southwest. This translates into a welfare gain of around \$77.3 million. As expected, this exceeds the welfare change under the exogenous carbon price; furthermore, our confidence levels now allow us to rule out a negative welfare impact.⁴⁹

6. CONCLUSIONS

We have developed GLC—a new, simple, flexible reduced-form model of imperfect competition that nests many existing oligopoly models as special cases. We showed how firm-level cost pass-through alone is a sufficient statistic for the profit impact of regulation on individual firms. Compared with existing literature, GLC relies heavily on supply linearity but allows near-arbitrary firm heterogeneity and does not require information about conduct parameters and markups or assumptions about rationality or equilibrium. We also showed how the GLC structure can be used for welfare analysis and to endogenize the extent of regulation. To illustrate GLC’s empirical usefulness, we estimated *ex ante* the impacts of future carbon pricing for US airlines. We found significant intra-industry heterogeneity in pass-through between low-cost and legacy carriers, driven by differences in product portfolios, cost structures and consumer demand. From a policy perspective, we therefore expect these carrier types to have very different incentives to embrace climate regulation.

We developed GLC with the objective of a simple yet robust modelling approach that can be put to use on a range of policy issues across different industries. We hope that GLC will prove useful in other single-industry contexts that, like airlines, are characterized by complex firm heterogeneity in demand, costs and conduct. This is common feature of many applications in industrial organization but it is also increasingly relevant for work in public economics, for

⁴⁹We can also decompose the “shortfall” in carbon pricing relative to the Pigouvian benchmark: the Buchanan market-power effect accounts for almost 90% of the shortfall—and is therefore empirically much stronger than the Grossman-Helpman lobbying effect. A detailed description is given in Appendix B.8.

example, studying the welfare impacts of a soda tax. GLC also lends itself to cross-industry analysis that seeks to apply a consistent economic structure across many different industries. This a common feature of work in macroeconomics and international trade, such as examining the impact of tariff reforms in reducing production costs, and in labour economics, such as studying the impact of a minimum wage. GLC’s comparative advantage lies in lower complexity in conducting *ex ante* policy evaluations, rooted in the theory of imperfect competition, but without requiring commitment to a particular model or notion of equilibrium.

REFERENCES

- [1] Acemoglu, Daron and Martin Kaare Jensen (2013). Aggregative Comparative Statics. *Games & Economic Behavior* 81, 27–49.
- [2] Allaz, Blaise and Vila, Jean-Luc (1993). Cournot Competition, Forward Markets and Efficiency. *Journal of Economic Theory* 59, 1-16.
- [3] 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.
- [4] Anderson, Simon P., André de Palma and Brent Kreider (2001). Tax Incidence in Differentiated Product Oligopoly. *Journal of Public Economics* 81, 173–192.
- [5] Armstrong, Mark and John Vickers (2018). Multiproduct Pricing Made Simple. *Journal of Political Economy* 126(4), 1444-1471.
- [6] Ashenfelter, Orley and Robert S. Smith (1979). Compliance with the Minimum Wage Law. *Journal of Political Economy* 87, 333–350.
- [7] 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.
- [8] 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.
- [9] Azar, José, Martin C. Schmalz and Isabel Tecu (2018). Anti-Competitive Effects of Common Ownership. *Journal of Finance*, forthcoming.
- [10] Bagnoli, Mark and Ted Bergstrom (2005). Log-concave Probability and its Applications. *Economic Theory* 26, 445–469.
- [11] Ballester, Coralio, Antoni Calvo-Armengol and Yves Zenou (2006). Who’s Who in Networks. Wanted: The Key Player. *Econometrica* 74(5), 1403–1417.
- [12] Baumol, William J. and Wallace E. Oates (1988). *The Theory of Environmental Policy*. Cambridge University Press, Cambridge: UK (2nd edition).
- [13] Bernheim, B. Douglas and Michael D. Whinston (1986). Menu Auctions, Resource Allocation and Economic Influence. *Quarterly Journal of Economics* 101, 1–31.
- [14] Berry, Steven, James Levinsohn and Ariel Pakes (1995). Automobile Prices in Market Equilibrium. *Econometrica* 63(4) 841–890.
- [15] Berry, Steven, and Panle Jia (2010). Tracing the Woes: An Empirical Analysis of the Airline Industry. *American Economic Journal: Microeconomics* 2(3): 1–43.

- [16] Borenstein, Severin (1992). The Evolution of US Airline Competition. *Journal of Economic Perspectives* 6, 45–73.
- [17] Borenstein, Severin (2011). Why Can't US Airlines Make Money? *American Economic Review* 101, 233–237.
- [18] Borenstein, Severin and Nancy L. Rose (1994). Competition and Price Dispersion in the US Airline Industry. *Journal of Political Economy* 102, 653–683.
- [19] 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.
- [20] Bramoullé, Yann, Rachel Kranton and Martin D'Amours (2014). Strategic Interaction and Networks. *American Economic Review* 104(3), 898–930.
- [21] Brander, James A. and Anming Zhang (1990). Market Conduct in the Airline Industry: An Empirical Investigation. *RAND Journal of Economics* 21, 567–583.
- [22] 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.
- [23] Buchanan, James (1969). External Diseconomies, Corrective Taxes, and Market Structure. *American Economic Review* 59, 174–177.
- [24] Bulow, Jeremy I. and Paul Pfleiderer (1983). A Note on the Effect of Cost Changes on Prices. *Journal of Political Economy* 91, 182–185.
- [25] 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.
- [26] Bushnell, James B., Erin T. Mansur, and Celeste Saravia (2008). Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets. *American Economic Review* 98(1), 237–266.
- [27] 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.
- [28] Busse, Meghan (2002). Firm Financial Condition and Airline Price Wars. *RAND Journal of Economics* 33, 298–318.
- [29] Chetty, Raj (2009). Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods, *Annual Review of Economics* 1, 451–487.
- [30] Ciliberto, Federico and Elie Tamer (2009). Market Structure and Multiple Equilibria in Airline Markets. *Econometrica* 77, 1791–1828.
- [31] Corchón, Luis C. (1994). Comparative Statics for Aggregative Games: The Strong Concavity Case. *Mathematical Social Sciences* 28, 151–165.
- [32] 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.
- [33] 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.
- [34] Daughety, Andrew (1990). Beneficial Concentration. *American Economic Review* 80, 1231–1237.

- [35] De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal and Nina Pavcnik (2016). Prices, Markups and Trade Reform. *Econometrica* 84, 445–510.
- [36] DellaVigna, Stefano and Matthew Gentzkow (2019). Uniform Pricing in U.S. Retail Chains. *Quarterly Journal of Economics* 134(4), 2011–2084.
- [37] Draca, Mirko, Stephen Machin and John Van Reenen (2011). Minimum Wages and Firm Profitability. *American Economic Journal: Applied Economics* 3, 129–151.
- [38] Einav, Liran and Jonathan Levin (2010). Empirical Industrial Organization: A Progress Report. *Journal of Economic Perspectives* 24, 145–162.
- [39] Ellison, Glen (2006). Bounded Rationality in Industrial Organization, in Richard Blundell, Whitney Newey, and Torsten Persson (eds.) *Advances in Economics and Econometrics: Theory and Applications*, Ninth World Congress, Cambridge University Press.
- [40] Elliott, Matthew and Andrea Galeotti (2019). The Role of Networks in Antitrust Investigations. *Oxford Review of Economic Policy* 35(4), 614–637.
- [41] European Environment Agency (2016). EEA Air Pollutant Emission Inventory Guidebook 2016, Chapter 1.A.3, Annex 4. Technical Report, 2016.
- [42] Fabra, Natalia and Mar Reguant (2014). Pass-Through of Emissions Costs in Electricity Markets. *American Economic Review* 104, 2972–2899.
- [43] 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.
- [44] Fowlie, Meredith, Mar Reguant and Stephen P. Ryan (2016). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy* 124, 249–302.
- [45] 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.
- [46] Gaudenzi, Barbara, and Alessandro Buccioli (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.
- [47] 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.
- [48] Genakos, Christos and Mario Pagliero (2019). Competition and Pass-Through: Evidence from Isolated Markets, CEPR DP13882.
- [49] Genesove, David and Wallace P. Mullin (1998). Testing Static Oligopoly Models: Conduct and Cost in the Sugar Industry, 1890–1914. *RAND Journal of Economics* 29(2), 355–377.
- [50] 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.
- [51] Goldberg, Pinelopi K. and Giovanni Maggi (1999). Protection for Sale: An Empirical Investigation. *American Economic Review* 89(5), 1135–1155.
- [52] 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.
- [53] Grossman, Gene M. and Elhanan Helpman (1994). Protection for Sale. *American Economic Review* 84, 833–850.

- [54] Hackner, Jonas (2000). A Note on Price and Quantity Competition in Differentiated Oligopolies. *Journal of Economic Theory* 93, 233–239.
- [55] Hazledine, Tim (2006). Price Discrimination in Cournot-Nash Oligopoly. *Economics Letters* 93, 413–420.
- [56] 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.
- [57] Heskett, James L., and W. Earl Sasser Jr. (2010). Southwest Airlines: In a Different World. Harvard Business School Case 910-419.
- [58] IAWG (2016). Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866. Interagency Working Group on Social Cost of Greenhouse Gases, United States Government, August 2016.
- [59] ICAO (2016). *2016 Environmental Report*. International Civil Aviation Organisation, Geneva, Switzerland.
- [60] Johnson, Justin P. and David P. Myatt (2003). Multiproduct Quality Competition: Fighting Brands and Product Line Pruning. *American Economic Review* 93, 748–774.
- [61] Johnson, Justin P. and David P. Myatt (2006). Multiproduct Cournot Oligopoly. *RAND Journal of Economics* 37, 583–601.
- [62] Johnson, Justin P. and Andrew Rhodes (2018). Multiproduct Mergers and Quality Competition. Working paper at Cornell and Toulouse, November 2018.
- [63] Kim, E. Han and Vijay Singal (1993). Mergers and Market Power: Evidence from the Airline Industry. *American Economic Review* 83, 549–569.
- [64] Klemperer, Paul D. and Margaret A. Meyer (1989). Supply Function Equilibria in Oligopoly under Uncertainty. *Econometrica* 57, 1243–1277.
- [65] Knittel, Christopher R. and Konstantinos Metaxoglou (2014). Estimation of Random Coefficient Demand Models: Two Empiricists’ Perspective. *Review of Economics and Statistics* 96(1), 34–59.
- [66] Kutlu, Levent (2017) A Conduct Parameter Model of Price Discrimination. *Scottish Journal of Political Economy* 64(5), 530–536.
- [67] Mankiw, N. Gregory and Matthew Weinzierl (2006). Dynamic Scoring: A Back-of-the-Envelope Guide. *Journal of Public Economics* 90, 1415–1433.
- [68] McCalman, Phillip (2004). Protection for Sale and Trade Liberalization: an Empirical Investigation. *Review of International Economics* 12(1), 81–94.
- [69] 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.
- [70] Menezes, Flavio M. and John Quiggin (2012). More Competitors or More Competition? Market Concentration and the Intensity of Competition. *Economics Letters* 117, 712–714.
- [71] Miklos-Thal, Jeanine and Greg Shaffer (2020). Pass-Through as an Economic Tool - On Exogenous Competition, Social Incidence, and Price Discrimination. *Journal of Political Economy*, forthcoming.
- [72] 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.
- [73] Miller, Nathan H., Marc Remer, Conor Ryan and Gloria Sheu (2016). Pass-Through and the Prediction of Merger Price Effects. *Journal of Industrial Economics* 64(4), 683–709.

- [74] 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.
- [75] Monderer, Dov and Lloyd Shapley (1996). Potential Games. *Games & Economic Behavior* 14, 124–143.
- [76] 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.
- [77] Muehlegger, Erich and Richard L. Sweeney (2020). Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom. *Review of Economics and Statistics*, forthcoming.
- [78] Nevo, Aviv (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica* 69, 307–342.
- [79] Nocke, Volker and Nicolas Schutz (2018). Multiproduct-Firm Oligopoly: An Aggregative Games Approach. *Econometrica* 86(2), 523–557.
- [80] Nordhaus, William D. (2017). Revisiting the Social Cost of Carbon. *Proceedings of the National Academy of Sciences* 114, 1518–1523.
- [81] 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.
- [82] 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.
- [83] Peltzman, Sam (2000). Prices Rise Faster Than They Fall. *Journal of Political Economy* 108(3), 466–502.
- [84] Pesaran, M. Hashem, and Ron P. Smith (1995). Estimating Long-Run Relationships from Dynamic Heterogenous Panels. *Journal of Econometrics* 69, 79–113.
- [85] Reiss, Peter, and Frank Wolak (2007). Structural Econometric Modeling: Rationales and Examples from Industrial Organisation. In *Handbook of Econometrics* Volume 6A.
- [86] 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.
- [87] Ritz, Robert A. (2014). On Welfare Losses Due to Imperfect Competition. *Journal of Industrial Economics* 62, 167–190.
- [88] Salop, Steven C. (1979). Monopolistic Competition with Outside Goods. *Bell Journal of Economics* 10:1, 141–156.
- [89] Seade, Jesus (1985). Profitable Cost Increases and the Shifting of Taxation: Equilibrium Response of Markets in Oligopoly. *Warwick Economics Research Paper Series* 260.
- [90] Stolper, Samuel (2016). Who Bears the Burden of Energy Taxes? The Role of Local Pass-Through. Harvard Environmental Economics Program Discussion Paper 16-70.
- [91] Spiegler, Ran (2011). *Bounded Rationality and Industrial Organization*, Oxford University Press.
- [92] Sutton, John (2007). Market Structure: Theory and Evidence. In: Mark Armstrong and Robert Porter (eds.), *Handbook of Industrial Organization*, Volume 3, 2301–2368.
- [93] 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.

- [94] Tully, Shawn (2015) Southwest Bets Big on Business Travelers. *Fortune*, September 23.
- [95] Varian, Hal R. (1985). Price Discrimination and Social Welfare. *American Economic Review* 75(4), 870–875.
- [96] Verboven, Frank (2012). *Empirical Industrial Organization and Competition Policy*, Unpublished book manuscript.
- [97] 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.
- [98] Zhang, Anming and Yahua Zhang (2017). Airline Economics: An Introductory Survey. Working Paper at University of British Columbia.

Appendices

A. PROOFS

Proof of Proposition 1. We begin with some preliminaries and then prove parts (a)–(c) of the result. In general, using A1, firm i 's profits as a function of regulation τ are given by $\Pi_i(\tau) = p_i(\tau)x_i(\tau) - C_i(x_i(\tau), e_i(\tau)) - \tau e_i(\tau)$. Using A2 and A3, it follows that $[C_i(x_i(\tau), e_i(\tau)) + \tau e_i(\tau)] = k_i(\tau)x_i(\tau)$ as $e_i(\tau) = z_i(\tau)x_i(\tau)$ and $k_i(\tau) = c_i(\tau) + \tau z_i(\tau)$, so firm i 's (optimized) profits can be written as $\Pi_i(\tau) = [p_i(\tau) - k_i(\tau)]x_i(\tau)$. We now pin down the change in firm i 's profits $\Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \equiv [\Pi_i(\bar{\tau}; \Phi) - \Pi_i(\underline{\tau}; \Phi)]$ in response to a regulation with scope Φ tightening from $\underline{\tau}$ to $\bar{\tau}$ under the respective assumptions of parts (a)–(c).

For part (a), a “small” tightening of regulation from $\underline{\tau}$ to $\bar{\tau}$ (with $\bar{\tau} \rightarrow \underline{\tau}$) means that the profit impact can be written as a linearization around the initial profit level:

$$\Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi)|_{\bar{\tau} \rightarrow \underline{\tau}} \simeq (\bar{\tau} - \underline{\tau}) \left. \frac{d\Pi_i(\tau; \Phi)}{d\tau} \right|_{\tau=\underline{\tau}}.$$

By A4, $x_i(\tau) = \psi_i[p_i(\tau) - k_i(\tau)]$, so firm i 's profits can be written as $\Pi_i(\tau) = \psi_i[p_i(\tau) - k_i(\tau)]^2$. Hence performing the differentiation gives:

$$\begin{aligned} \left. \frac{d\Pi_i(\tau; \Phi)}{d\tau} \right|_{\tau=\underline{\tau}} &= 2\psi_i[p_i(\underline{\tau}) - k_i(\underline{\tau})] \left(\left. \frac{dp_i(\tau)}{d\tau} \right|_{\tau=\underline{\tau}} - \left. \frac{dk_i(\tau)}{d\tau} \right|_{\tau=\underline{\tau}} \right) \\ &= 2x_i(\underline{\tau}) \left(\left. \frac{dp_i(\tau)}{d\tau} \right|_{\tau=\underline{\tau}} - \left. \frac{dk_i(\tau)}{d\tau} \right|_{\tau=\underline{\tau}} \right) \\ &= 2x_i(\underline{\tau})[\rho_i(\underline{\tau}; \Phi) - 1]z_i(\underline{\tau}) \\ &= -2e_i(\underline{\tau})[1 - \rho_i(\underline{\tau}; \Phi)], \end{aligned}$$

where the second equality again uses A4, the third equality uses $dk_i(\tau)/d\tau = z_i(\tau)$, which follows from A1–A3 and the envelope theorem, as well as the definition of the marginal rate of firm-level cost pass-through $\rho_i(\tau; \Phi) \equiv [dp_i(\tau; \Phi)/d\tau]/[dk_i(\tau)/d\tau]$, and the fourth equality uses the definition of the regulatory intensity $z_i \equiv e_i/x_i$. This completes the proof of part (a).

For part (b), the change in firm i 's profits in response to a “large” tightening in regulation from $\underline{\tau}$ to $\bar{\tau}$ can be written as the integral of small changes as characterized in part (a), $\Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) = 2 \int_{\tau=\underline{\tau}}^{\bar{\tau}} [d\Pi_i(\tau; \Phi)/d\tau] d\tau$. Alternatively it can be expressed as:

$$\begin{aligned} \Delta\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) &= [p_i(\underline{\tau}) - k_i(\underline{\tau})]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) + [\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]x_i(\underline{\tau}) + \\ &\quad [\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi), \end{aligned}$$

where note that the unit cost increase $\Delta k_i(\bar{\tau}, \underline{\tau})$ does not depend on Φ as it is internal to firm i (by A1–A3). The first two terms together make up the “first-order effect” on profits and the third term represents the “second-order effect”. Next we derive expressions for these two effects in terms of pass-through.

For the first-order effect, by A4, $x_i(\tau) = \psi_i[p_i(\tau) - k_i(\tau)]$ so that the change in firm i 's output in response to tighter regulation satisfies:

$$\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) = \psi_i[\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})].$$

Using this, the two parts of the first-order effect can be re-written in combined form as:

$$[p_i(\underline{\tau}) - k_i(\underline{\tau})]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) + [\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]x_i(\underline{\tau}) = 2x_i(\underline{\tau})[\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]$$

Next we rewrite this in terms of firm i 's initial use of the regulated factor as follows:

$$\begin{aligned}
2x_i(\underline{\tau})[\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})] &= 2x_i(\underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]\Delta k_i(\bar{\tau}, \underline{\tau}) \\
&= 2e_i(\underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]\frac{\Delta k_i(\bar{\tau}, \underline{\tau})}{z_i(\underline{\tau})} \\
&= 2e_i(\underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]\frac{\int_{\tau=\underline{\tau}}^{\bar{\tau}} z_i(\tau)}{z_i(\underline{\tau})} \\
&= 2(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]g_i(\bar{\tau}, \underline{\tau}),
\end{aligned}$$

where the first equality uses the definition of the average rate of firm-level cost pass-through, $\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) \equiv \Delta p_i(\bar{\tau}, \underline{\tau}; \Phi)/\Delta k_i(\bar{\tau}, \underline{\tau})$, the second equality uses the definition of the regulatory intensity $z_i \equiv e_i/x_i$, the third equality uses $\Delta k_i(\bar{\tau}, \underline{\tau}) = \int_{\tau=\underline{\tau}}^{\bar{\tau}} z_i(\tau)$, which follows from A1–A3 and the envelope theorem, and the fourth equality uses our definition of the extent of factor substitution, $g_i(\bar{\tau}, \underline{\tau}) = [\int_{\tau=\underline{\tau}}^{\bar{\tau}} z_i(\tau)d\tau]/[(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})] > 0$.

For the second-order effect, similarly, we rewrite as follows:

$$\begin{aligned}
[\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) &= \Delta k_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) \\
&= (\bar{\tau} - \underline{\tau})z_i(\underline{\tau})g_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) \\
&= (\bar{\tau} - \underline{\tau})e_i(\underline{\tau})g_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]\frac{\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi)}{x_i(\underline{\tau})},
\end{aligned}$$

where the first equality again uses the definition of average pass-through $\bar{\rho}_i$, the second equality uses the relationship $\Delta k_i(\bar{\tau}, \underline{\tau}) = \int_{\tau=\underline{\tau}}^{\bar{\tau}} z_i(\tau)d\tau = (\bar{\tau} - \underline{\tau})z_i(\underline{\tau})g_i(\bar{\tau}, \underline{\tau})$, and the third equality uses the definition of the regulatory intensity $z_i \equiv e_i/x_i$. The output change follows as:

$$\begin{aligned}
\frac{\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi)}{x_i(\underline{\tau})} &= \frac{\psi_i[\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]}{\psi_i[p_i(\underline{\tau}) - k_i(\underline{\tau})]} \\
&= \frac{\Delta k_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_i(\underline{\tau}) - k_i(\underline{\tau})]} \\
&= \frac{(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})g_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_i(\underline{\tau}) - k_i(\underline{\tau})]} \\
&= \frac{[(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})/p_i(\underline{\tau})]g_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_i(\underline{\tau}) - k_i(\underline{\tau})]/p_i(\underline{\tau})},
\end{aligned}$$

where the first equality uses A4, the second equality uses the definition of pass-through $\bar{\rho}_i$, the third uses the definition of the extent of factor substitution g_i , and the fourth equality is simple rearranging. This finally allows us to write the second-order effect as:

$$\begin{aligned}
[\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) &= (\bar{\tau} - \underline{\tau})e_i(\underline{\tau})g_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1] \times \\
&\quad \times \frac{[(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})/p_i(\underline{\tau})]g_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_i(\underline{\tau}) - k_i(\underline{\tau})]/p_i(\underline{\tau})}.
\end{aligned}$$

Using these expressions for the first- and second-order effects in the initial expression for the

profit impact of firm i shows that:

$$\begin{aligned}
\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) &= [p_i(\underline{\tau}) - k_i(\underline{\tau})]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) + [\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]x_i(\underline{\tau}) + \\
&\quad + [\Delta p_i(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_i(\bar{\tau}, \underline{\tau})]\Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) \\
&= 2[(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})][\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]g_i(\bar{\tau}, \underline{\tau}) \left(1 + \frac{[(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})/p_i(\underline{\tau})]g_i(\bar{\tau}, \underline{\tau})[\bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{2[p_i(\underline{\tau}) - k_i(\underline{\tau})]/p_i(\underline{\tau})} \right) \\
&= -2[(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})][1 - \bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi)]g_i(\bar{\tau}, \underline{\tau}) \left(1 - \frac{g_i(\bar{\tau}, \underline{\tau})[1 - \bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})/p_i(\underline{\tau})]}{2[p_i(\underline{\tau}) - k_i(\underline{\tau})]/p_i(\underline{\tau})} \right) \\
&= -2[1 - \bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})]\Omega_i(\bar{\tau}, \underline{\tau}; \Phi),
\end{aligned}$$

where the final equality uses the definition

$$\Omega_i(\bar{\tau}, \underline{\tau}; \Phi) \equiv g_i(\bar{\tau}, \underline{\tau}) \left(1 - \frac{g_i(\bar{\tau}, \underline{\tau})[1 - \bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})/p_i(\underline{\tau})]}{2[p_i(\underline{\tau}) - k_i(\underline{\tau})]/p_i(\underline{\tau})} \right)$$

as claimed, and $\Omega_i(\bar{\tau}, \underline{\tau}; \Phi) > 0$ follows because $x_i(\underline{\tau}) + \Delta x_i(\bar{\tau}, \underline{\tau}; \Phi) > 0$, that is, firm i remains active, by assumption (as is implicit in A4) following the tightening of regulation. This completes part (b) of the result.

For part (c), under the two maintained assumptions (i) cost increase is modest relative to its initial price, that is, $[(\bar{\tau} - \underline{\tau})z_i(\underline{\tau})/p_i(\underline{\tau})]$ is “small”, and (ii) regulatory intensity is approximately constant, that is, $z_i(\bar{\tau}) \simeq z_i(\underline{\tau})$, the result follows immediately from the expression derived in part (b). In particular, the second assumption implies $g_i(\bar{\tau}, \underline{\tau}) \simeq 1$ which, combined with the first assumption, then implies $\Omega_i(\bar{\tau}, \underline{\tau}; \Phi) \simeq 1$, and so:

$$\Delta \Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \simeq -2[1 - \bar{\rho}_i(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})],$$

as claimed, thus completing the final part (c) of the proof.

Proof of Proposition 2. In general, using A1M, firm i 's profits across its offering as a function of regulation τ is given by: $\Pi_i = \sum_{m \in M} p_{im}x_{im} - C_i(\mathbf{x}_i, \mathbf{e}_i) - \tau e_i$. Using multidimensional GLC's A2M and A3M, it follows that $[C_i(\mathbf{x}_i(\tau), \mathbf{e}_i(\tau)) + \tau e_i(\tau)] = \sum_{m \in M} k_{im}(\tau)x_{im}(\tau)$ as $e_{im}(\tau) = z_{im}(\tau)x_{im}(\tau)$ and $k_{im}(\tau) = c_{im}(\tau) + \tau z_{im}(\tau)$, so firm i 's (optimized) profits become:

$$\Pi_i(\tau) = \sum_{m \in M} \{ [p_{im}(\tau) - k_{im}(\tau)] x_{im}(\tau) \}.$$

We now prove parts (a)–(c) of the result. For part (a), by A4M, $x_{im}(\tau) = \psi_{im}[p_{im}(\tau) - k_{im}(\tau)]$, and so firm i 's profits can be written as $\Pi_i(\tau) = \sum_{m \in M} \psi_{im}[p_{im}(\tau) - k_{im}(\tau)]^2$. With a “small” tightening of a regulation τ with scope Φ from $\underline{\tau}$ to $\bar{\tau}$ (with $\bar{\tau} \rightarrow \underline{\tau}$), the profit impact can be written as a linearization around the initial profit level:

$$\Delta \Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \Big|_{\bar{\tau} \rightarrow \underline{\tau}} \simeq (\bar{\tau} - \underline{\tau}) \frac{d\Pi_i(\tau; \Phi)}{d\tau} \Big|_{\tau = \underline{\tau}}.$$

Hence performing the differentiation gives:

$$\begin{aligned}
\frac{d\Pi_i(\tau; \Phi)}{d\tau} \Big|_{\tau = \underline{\tau}} &= 2 \sum_{m \in M} x_{im}(\underline{\tau}) \left(\frac{dp_{im}(\tau)}{d\tau} \Big|_{\tau = \underline{\tau}} - \frac{dk_{im}(\tau)}{d\tau} \Big|_{\tau = \underline{\tau}} \right) \\
&= 2 \sum_{m \in M} x_{im}(\underline{\tau}) [\rho_{im}(\underline{\tau}; \Phi) - 1] z_{im}(\underline{\tau}) \\
&= 2 \sum_{m \in M} e_{im}(\underline{\tau}) [\rho_{im}(\underline{\tau}; \Phi) - 1] \\
&= -2e_i(\underline{\tau}) [1 - \rho_i^M(\underline{\tau}; \Phi)],
\end{aligned}$$

where the first equality uses A4M, the second equality uses the definition of marginal cost pass-through for component $m \in M$, $\rho_{im}(\tau; \Phi) \equiv [dp_{im}(\tau; \Phi)/d\tau]/[dk_{im}(\tau)/d\tau]$ as well as $dk_{im}(\tau)/d\tau = z_{im}(\tau)$ from the envelope theorem (A1M-A3M), the third equality uses definition of the regulatory intensity $z_{im} \equiv e_{im}/x_{im}$ for component $m \in M$, and the fourth equality uses the definition of multidimensional cost pass-through, $\rho_i^M(\tau; \Phi) \equiv \sum_{m \in M} \omega_{im}(\tau) \rho_{im}(\tau; \Phi)$, where $\omega_{im}(\tau) \equiv e_{im}(\tau)/e_i(\tau)$ as well as $\sum_{m \in M} \omega_{im}(\tau) = 1$.

For part (b), the change in firm i 's profits in response to a “large” tightening in regulation from $\underline{\tau}$ to $\bar{\tau}$ can be written as the integral of small changes as characterized in part (a), $\Delta \Pi_i(\bar{\tau}, \underline{\tau}; \Phi) = 2 \int_{\tau=\underline{\tau}}^{\bar{\tau}} [d\Pi_i(\tau; \Phi)/d\tau] d\tau$. Alternatively it can be expressed as:

$$\Delta \Pi_i(\bar{\tau}, \underline{\tau}; \Phi) = \sum_{m \in M} \left\{ \begin{aligned} & [p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})] \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) + [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})] x_{im}(\underline{\tau}) + \\ & + [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})] \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) \end{aligned} \right\}$$

where note that the unit cost increases $\Delta k_{im}(\bar{\tau}, \underline{\tau})$ do not depend on Φ as they are internal to firm i (by A1M-A3M). The first two terms together make up the “first-order effect” on profits and the third term represents the “second-order effect”. Next we derive expressions for these two effects in terms of pass-through.

For the first-order effect, by A4M, $x_{im}(\tau) = \psi_{im}[p_{im}(\tau) - k_{im}(\tau)]$ so that the change in firm i 's output of component m of its offering in response to tighter regulation satisfies:

$$\Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) = \psi_{im}[\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})].$$

Using this, the two parts of the first-order effect can be re-written in combined form as:

$$\begin{aligned} & \sum_{m \in M} \{ [p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})] \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) + [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})] x_{im}(\underline{\tau}) \} \\ & = 2 \sum_{m \in M} x_{im}(\underline{\tau}) [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})] \end{aligned}$$

Next we rewrite this in terms of firm i 's initial use of the regulated factor as follows:

$$\begin{aligned} 2 \sum_{m \in M} x_{im}(\underline{\tau}) [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})] & = 2 \sum_{m \in M} x_{im}(\underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] \Delta k_{im}(\bar{\tau}, \underline{\tau}) \\ & = 2 \sum_{m \in M} e_{im}(\underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] \frac{\Delta k_{im}(\bar{\tau}, \underline{\tau})}{z_{im}(\underline{\tau})} \\ & = 2 \sum_{m \in M} e_{im}(\underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] \frac{\int_{\tau=\underline{\tau}}^{\bar{\tau}} z_{im}(\tau) d\tau}{z_{im}(\underline{\tau})} \\ & = 2(\bar{\tau} - \underline{\tau}) \sum_{m \in M} e_{im}(\underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] g_{im}(\bar{\tau}, \underline{\tau}), \end{aligned}$$

where the first equality uses the definition of the average rate of firm-level cost pass-through on component m , $\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) \equiv \Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi)/\Delta k_{im}(\bar{\tau}, \underline{\tau})$, the second equality uses the definition of the regulatory intensities $z_{im} \equiv e_{im}/x_{im}$, the third equality uses $\Delta k_{im}(\bar{\tau}, \underline{\tau}) = \int_{\tau=\underline{\tau}}^{\bar{\tau}} z_{im}(\tau) d\tau$, which follows from A1M-A3M and the envelope theorem, and the fourth equality uses our definition of the extent of factor substitution, $g_{im}(\bar{\tau}, \underline{\tau}) = [\int_{\tau=\underline{\tau}}^{\bar{\tau}} z_{im}(\tau) d\tau]/[(\bar{\tau} - \underline{\tau}) z_{im}(\underline{\tau})] > 0$.

For the second-order effect, similarly, we rewrite as follows:

$$\begin{aligned}
\sum_{m \in M} [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})] \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) &= \sum_{m \in M} \Delta k_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) \\
&= (\bar{\tau} - \underline{\tau}) \sum_{m \in M} z_{im}(\underline{\tau}) g_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) \\
&= (\bar{\tau} - \underline{\tau}) \sum_{m \in M} e_{im}(\underline{\tau}) g_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] \frac{\Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi)}{x_{im}(\underline{\tau})},
\end{aligned}$$

where the first equality again uses the definition of average pass-through $\bar{\rho}_{im}$, the second equality uses the relationship $\Delta k_{im}(\bar{\tau}, \underline{\tau}) = \int_{\underline{\tau}}^{\bar{\tau}} z_{im}(\tau) d\tau = (\bar{\tau} - \underline{\tau}) z_{im}(\underline{\tau}) g_{im}(\bar{\tau}, \underline{\tau})$, and the third equality uses the definition of the regulatory intensity $z_{im} \equiv e_{im}/x_{im}$. The output change follows as:

$$\begin{aligned}
\frac{\Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi)}{x_{im}(\underline{\tau})} &= \frac{\psi_{im} [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})]}{\psi_{im} [p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})]} \\
&= \frac{\Delta k_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})]} \\
&= \frac{(\bar{\tau} - \underline{\tau}) z_{im}(\underline{\tau}) g_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})]} \\
&= \frac{[(\bar{\tau} - \underline{\tau}) z_{im}(\underline{\tau}) / p_{im}(\underline{\tau})] g_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})] / p_{im}(\underline{\tau})},
\end{aligned}$$

where the first equality uses A4M, the second equality uses the definition of pass-through $\bar{\rho}_{im}$, the third uses the definition of the extent of factor substitution g_{im} , and the fourth equality is simple rearranging. This finally allows us to write the second-order effect as:

$$\begin{aligned}
&\sum_{m \in M} [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})] \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) \\
&= (\bar{\tau} - \underline{\tau}) \sum_{m \in M} \left\{ e_{im}(\underline{\tau}) g_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] \frac{[(\bar{\tau} - \underline{\tau}) z_{im}(\underline{\tau}) / p_{im}(\underline{\tau})] g_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})] / p_{im}(\underline{\tau})} \right\}.
\end{aligned}$$

Using these expressions for the first- and second-order effects in the initial expression for the profit impact of firm i shows that:

$$\begin{aligned}
\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) &= \sum_{m \in M} \left\{ [p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})] \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) + [\Delta p_{im}(\bar{\tau}, \underline{\tau}; \Phi) - \Delta k_{im}(\bar{\tau}, \underline{\tau})] x_{im}(\underline{\tau}) \right\} \\
&= 2(\bar{\tau} - \underline{\tau}) \sum_{m \in M} e_{im}(\underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] g_{im}(\bar{\tau}, \underline{\tau}) + \\
&+ (\bar{\tau} - \underline{\tau}) \sum_{m \in M} \left\{ e_{im}(\underline{\tau}) g_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1] \frac{[(\bar{\tau} - \underline{\tau}) z_{im}(\underline{\tau}) / p_{im}(\underline{\tau})] g_{im}(\bar{\tau}, \underline{\tau}) [\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi) - 1]}{[p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})] / p_{im}(\underline{\tau})} \right\} \\
&= -2(\bar{\tau} - \underline{\tau}) \left[\sum_{m \in M} e_{im}(\underline{\tau}) [1 - \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)] g_{im}(\bar{\tau}, \underline{\tau}) - \right. \\
&\quad \left. \sum_{m \in M} \left\{ e_{im}(\underline{\tau}) g_{im}(\bar{\tau}, \underline{\tau}) [1 - \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)] \frac{[(\bar{\tau} - \underline{\tau}) z_{im}(\underline{\tau}) / p_{im}(\underline{\tau})] g_{im}(\bar{\tau}, \underline{\tau}) [1 - \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)]}{2[p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})] / p_{im}(\underline{\tau})} \right\} \right] \\
&= -2 \sum_{m \in M} [1 - \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)] [(\bar{\tau} - \underline{\tau}) e_{im}(\underline{\tau})] \Omega_{im}(\bar{\tau}, \underline{\tau}; \Phi),
\end{aligned}$$

where the final equality uses the definition

$$\Omega_{im}(\bar{\tau}, \underline{\tau}; \Phi) \equiv g_{im}(\bar{\tau}, \underline{\tau}) \left(1 - \frac{g_{im}(\bar{\tau}, \underline{\tau})[1 - \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})z_{im}(\underline{\tau})/p_{im}(\underline{\tau})]}{2[p_{im}(\underline{\tau}) - k_{im}(\underline{\tau})]/p_{im}(\underline{\tau})} \right)$$

as claimed, and $\Omega_{im}(\bar{\tau}, \underline{\tau}; \Phi) > 0$ follows because $x_{im}(\underline{\tau}) + \Delta x_{im}(\bar{\tau}, \underline{\tau}; \Phi) > 0$, that is, firm i remains active on each element m , by assumption (as is implicit in A4M) following the tightening of regulation. This completes part (b) of the result.

For part (c), under the two maintained assumptions that (i) the cost increase on each component m is modest relative to its initial price, that is, $[(\bar{\tau} - \underline{\tau})z_{im}(\underline{\tau})]/p_{im}(\underline{\tau})$ is “small” for all $m \in M$, and (ii) its regulatory intensity on each component m is approximately constant, that is, $z_{im}(\bar{\tau}) \simeq z_{im}(\underline{\tau})$, the result follows immediately from the expression derived in part (b). In particular, the second assumption implies $g_{im}(\bar{\tau}, \underline{\tau}) \simeq 1$ for all $m \in M$ which, combined with the first assumption, then implies $\Omega_{im}(\bar{\tau}, \underline{\tau}; \Phi) \simeq 1$ for all $m \in M$, and so:

$$\Pi_i(\bar{\tau}, \underline{\tau}; \Phi) \simeq -2 \sum_{m \in M} [1 - \bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})e_{im}(\underline{\tau})] = -2[1 - \bar{\rho}_i^M(\bar{\tau}, \underline{\tau}; \Phi)][(\bar{\tau} - \underline{\tau})e_i(\underline{\tau})]$$

using the definition of multidimensional average pass-through $\bar{\rho}_i^M(\bar{\tau}, \underline{\tau}; \Phi) \equiv \sum_{m \in M} \omega_{im}(\underline{\tau})\bar{\rho}_{im}(\bar{\tau}, \underline{\tau}; \Phi)$, and thus completing the final part (c) of the proof.

B. DATA CONSTRUCTION AND FURTHER RESULTS

B.1. Data construction. 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 2013Q4 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: $k_{ilj} = b_{ilt}^0 + b_{ilt}^1 d_{ilj}$. 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 $\frac{b_{ilt}^0}{b_{ilt}^1} = 131$ 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 , $\tilde{d}_{iljt} = (\frac{b_{ilt}^0}{b_{ilt}^1} + distance_j) \times dep_performed_{iljt}$, so that $\alpha_{ijt}(l) = \frac{\tilde{d}_{iljt}}{\sum_l \tilde{d}_{iljt}}$. Using these shares we construct the weighted average fuel cost $k_{ijt} = \sum_l \alpha_{ijt}(l)k_{ilt}$.

Non-fuel cost index 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 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 2004Q1 value for American Airlines.

GDP growth 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 2004 we interpolated the annual data as quarterly data are not available.

Competitors and low cost competitors (LCC) are constructed using the full 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. LLC competitors is the number of competitors from the set of low-cost carriers, as defined by ICAO.

Carrier size is constructed using the full DB1A. It is the sum of all passengers on all routes in a given quarter for a give airline.

B.2. Results for individual legacy carriers. Table B1 gives descriptive statistics for the individual carriers in our sample. These confirm that the legacy carriers form a group distinct from Southwest, for example on the basis of price, fuel cost or distance.

Table B2 contains the results of estimating a variant of Equation (4) with an interaction term for each individual legacy carrier, rather than a single dummy for all legacy carriers (Table B3 provides the results from the first stage estimates). It therefore provides the pass-through rate for each airline. Column 1 is the baseline result in the main text, and column 2 gives the pass-through results by airline. Southwest’s pass-through remains unchanged. The average of the individual legacy carriers’ pass-through rates is in line with the overall legacy carrier result in the baseline. There is, however, considerable heterogeneity among the legacy carriers. Delta has a particularly large pass-through rate. Despite this range, the results do support our prior (following the airlines literature) to consider Southwest as distinct from the legacy carriers: it is the only airline with a pass-through rate statistically significantly greater than 1.

B.3. Robustness. Table B4 contains the results with additional control variables that plausibly could impact airlines’ pricing decisions. Columns 1-5 show the impact of using alternative measures of demand. All but two of the demand controls are self-explanatory. Gravity demand (analogous to the concept as used in the trade literature) is a measure of demand on a route constructed as the product of (per capita GDP times population) in origin and destination cities, divided by distance. Network density is the number of actual connections (i.e. routes) between airports in a carrier’s network, divided by the total number of possible connections (the standard definition in the networks literature). Column 1 is the same baseline as in the main text, and contains the only measure of demand that is statistically significant. In whichever measure is used, the pass-through rates are very stable around the baseline result.

Column 6 includes a dummy for potential entry by Southwest, using the definition in Goolsbee and Syverson (2008). We don’t find a significant effect, and the pass-through rates are unchanged from the baseline. Column 7 reports the effect of bankruptcy (all carriers except for Southwest were bankrupt at some point in this period). This is significant, but does not statistically affect pass-through.

In addition to investigating the effect of controls on prices, we also interacted these variables with fuel costs, to see if they had an effect directly on pass-through rates. We found almost no results of statistical significance. We also experimented with varying the start and end dates of our time period, which again had little impact on our estimates. These results are omitted in the interests of space, but available on request.

B.4. Mean Group regressions. As a final robustness exercise, we allow for full heterogeneity across routes within a given carrier’s portfolio. The relationship of interest is a variant of equation (4) given by:

$$p_{ijt} = \rho_{ij}k_{ijt} + X'_{ijt}\beta_{ij} + \lambda_{ijt} + \eta_{ij} + \epsilon_{ijt}. \quad (9)$$

Carrier-specific pass-through rates ρ_i are obtained from equation (9) by running a separate regression for each carrier-route, and then taking a weighted average of the carrier-route ρ_{ij} to obtain a pass-through rate at the carrier level, ρ_i . The weights are emissions on the carrier-route. This approach imposes no homogeneity restrictions on the parameters across carrier-routes within an airline, which could be important given heterogeneities across routes in the airline industry. In running a separate regression for each product, we take a similar approach to Atkin and Donaldson (2015). The procedure could also be considered a special (non-dynamic) case of the "Mean Group" estimator in Pesaran and Smith (1995). 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.

A drawback of using this approach is that, unlike in the standard fixed-effects baseline, year-and-quarter time effects are not identified. Equation (9) is therefore estimated only with quarterly time effects. Instruments and first stage use the analogue of Equation (8). The results are given in Table B5. Southwest’s pass-through rate is 1.24, significantly above 1, and the legacy carriers’ pass-through rate is 0.75, significantly below 1, qualitatively and quantitatively very similar to the baseline estimates.

B.5. Results in support of GLC’s A4M. We conduct two empirical tests to provide evidence that demand does not seem to exhibit significant non-linearities. First, we estimate the pass-through in monopoly routes using our baseline model (7). Results are reported in Table B6. The estimated pass-through coefficient is 0.6 and not statistically different from 0.5 ($\Pr > \chi^2 = 0.66$), which is the pass-through predicted by a monopoly model with linear demand. This is very much in line with the results from Genakos and Pagliero (2019), who find a similar result for monopolists in retail petroleum markets. This result provides evidence that demand convexity does not seem to be significant in monopoly markets. The rest of the coefficients are in line with the results in the baseline model, which is reassuring.

Second, we directly test for the existence of non-linear demand terms in own price (GLC allows for non-linearities in rivals’ prices) in duopoly and triopoly markets.⁵⁰ We estimate a simple semi-linear demand model of the number of passengers on carrier-route ij in quarter t on own and rival prices, including squared terms, time effects and ij fixed effects. All prices are instrumented using the same instruments as in equation (8). Table B7 reports the results for the duopoly markets, whereas Table B8 for the triopoly. In duopoly markets (Table B7, column 1) we can see that the OLS own and cross price coefficients are significantly negative and positive, respectively, in line with expectations. Instrumenting both prices in column 2 moves the own price coefficient further away from zero. However, adding a squared term for own price in column 3 is not significant (and similarly when moving from column 5 to column 6). Remember, that GLC allows for non-linearities in rivals’ prices, so we just want to test whether a non-linear term on own price is improving the model fit. Likewise, in Table B8, column 1, own and cross price coefficients are significant and have the expected sign for a triopoly market. However, adding again a squared term for own price does not improve the fit of the model, as it is not statistically significant. Although the data for triopoly markets is considerably smaller and hence more noisy, none of the non-linear coefficients seem to be significant. Therefore, results seems to suggest that there are no significant demand non-linearities in this case, which together with assumptions A1M-A3M provide empirical support to also maintain assumption A4M.

B.6. Results for auxiliary assumptions of Proposition 2(c). We here discuss the evidence related to the two auxiliary assumptions for Proposition 2(c): (i) the cost shock due to regulation is "small" relative to product price and (ii) there is limited scope for factor substitution.

From Table 2, we can see that the emissions costs at a \$30/tCO₂ are about 3% of the average ticket price for both Southwest (emissions cost/price = $(4 \div 154.73) \times 100 = 2.59\%$) and legacy ((emissions cost/price = $(6.02 \div 227.75) \times 100 = 2.64\%$) carriers. We interpret this cost shock as being “small” relative to the average ticket price (also remember that the sample only includes domestic economy class tickets).

Regarding factor substitution, clearly both Southwest and legacy carriers have limited scope for large-scale jet fuel substitution. However, in practice, airlines can increase fuel efficiency by adjusting various aspects of a flight (e.g., route and altitude to minimize the impact of wind, holding patterns and taxiing between runway and terminal, etc.). We transformed fuel prices to carbon prices ($p_{carbon} = 0.00957 \times 30 \times p_{fuel}$) using the known constant technological relationship (also given in the main text) and the \$30/tCO₂ SCC estimate from IAWG (2016). We then estimated a log-log specification of carbon price on emissions, while using the same controls (GDP growth, carrier size, number of competitors) and fixed effects (time and carrier-route) as in our baseline specification. We estimate empirically that a 1% increase in carbon prices decreases the emissions’ intensity by -0.03% for legacy carriers to -0.01% for Southwest

⁵⁰We have an insufficient number of observations to run any regression for markets with four players or more.

(results not reported here, available on request). We again interpret these results as providing evidence for limited factor substitutability.

B.7. Implementation of Proposition 3. We implement our theoretical results from Proposition 3(a) for consumer welfare and 3(b) for social welfare using the following steps. First, we utilize the appropriate level of regulation (τ): In Table 4, column 3 we use an exogenous \$30/tCO₂, whereas in column 6 we use the endogenous \$17.68/tCO₂ (explained in the following subsection). Second, we use our estimates of airline firm-level pass-through to obtain an industry-average multidimensional pass-through rate of $\tilde{\rho}^M = 0.853$ (see Table 4) together with industry emissions in the last year of our sample ($E(\bar{x}) = 15,420,727$). Third, for the social cost of carbon (SCC) we set $D' = \$30/\text{tCO}_2$. Fourth, we derive $dE(\tau)/d\tau$ by estimating the tax elasticity of industry emissions $\eta(\tau) \equiv \frac{d \ln E(\tau)}{d \ln \tau}$. We calculated total carrier emissions at the route level and then estimated a log-log specification of carbon price on total emissions, while controlling for GDP growth, carrier size, number of competitors, quarter and route fixed effects (results not reported here, available on request). We obtain a statistically significant coefficient of $\eta = -0.253$ (*s.e.* = 0.017), in line with the prior literature (for example, the midpoint of the range of estimates in Fukui and Miyoshi (2017) is $\eta = -0.256$). As expected, airline emissions respond negatively to carbon pricing—but the elasticity is not large. The implied $\frac{dE(\tau)}{d\tau}$ can then be backed out using the elasticity formula.

B.8. Implementation of Proposition 4. As noted in the main text, the question of whether regulation can be treated as a “manna from heaven” arises naturally for an influential industry like US airlines. We therefore apply Proposition 4 using our firm-level pass-through estimates for airlines to obtain the political equilibrium carbon price for US airlines. More precisely, we consider a domestic US policymaker who chooses her utility-maximizing level of “complete regulation” for all airlines, cognizant of market power in airline markets and under influence of political lobbying by the airline industry.

We make our theoretical result from Proposition 4 operational in four steps. First, for the social cost of carbon (SCC) we again set $D' = \$30/\text{tCO}_2$. Second, for the lobbying parameter λ , we turn to the literature. Goldberg and Maggi (1999) were the first to empirically estimate this parameter, finding $\lambda = 0.02$ for the US. McCalman (2004) and Mitra and Ulubasoglu (2002) obtain similar results for Australia and Turkey respectively while Gawande and Bandyopadhyay (2000) find a much higher $\lambda = 0.5$. Based on these findings, we take $\lambda = 0.1$ as our baseline. Third, we use our estimates of firm-level pass-through for airlines to obtain an industry-average multidimensional pass-through rate of $\tilde{\rho}^M = 0.853$ (see Table 4). As industry-average pass-through lies below 100%, the industry collectively opposes a carbon tax and lobbies the government to water it down. Fourth, we use our estimate of $\eta = -0.253$.

Using these parameters in the formula from Proposition 4 yields an endogenous carbon price τ^* equal to \$17.68/tCO₂, which is 41% lower than the exogenous carbon price set equal to the SCC following the Pigouvian rule. This means that the profit impacts of regulation are less pronounced for an endogenous carbon price: this mitigates the losses experienced by the legacy carriers but conversely also limits the gains accruing to Southwest (see Table 4). This translates into a welfare gain of around \$77.3 million which, as expected, exceeds that under the exogenous carbon price (at \$30/tCO₂).

We can decompose the “shortfall” in carbon pricing relative to the Pigouvian benchmark into the two underlying distortions: market power and political lobbying.⁵¹ Observe that the Buchanan market-power effect is logically prior in that it can exist in principle without any Grossman-Helpman political lobbying (i.e., where $\lambda = 0$) but the reverse is not true.⁵² We therefore re-run the calculation with $\lambda = 0$, i.e. with the lobbying channel switched off, and find an endogenous carbon price of \$18.98. We conclude that the Buchanan effect is empirically much stronger than the Grossman-Helpman effect as it explains around 90% of the disparity between the endogenous and exogenous carbon prices.

⁵¹It is worth stressing that the shortfall in carbon pricing here is not driven by the “incompleteness” regulation where a carbon price applies only to a subset of firms competing in an industry (Fowlie et al. 2016).

⁵²In our setting, without any market power, there are no profits and hence nothing to lobby over and so the Grossman-Helpman effect is zero.

In sum, we can draw two main conclusions for US airlines. First, the political-equilibrium carbon price lies significantly below the SCC and, second, the shortfall in carbon pricing is primarily driven by market power with a small additional role played by lobbying. Hence this analysis may help explain why aviation has, so far, been a climate laggard. Looking ahead, it also suggests that policies to address market power may be able to complement policies to address environmental externalities.

We have re-run the numbers with different assumptions for the key parameters on political lobbying (λ) and the emissions elasticity (η). In particular, we consider a range of values $\lambda \in [0, \frac{1}{2}]$ based on the literature discussed above combined with a range $\eta \in [-.287, -.220]$ based on the 95% confidence intervals of our estimates and in accordance with prior literature. The resulting estimates for the endogenous carbon price τ^* vary within a range of \$13 – \$20, suggesting that our baseline estimate is quite robust.

The conclusions are relatively more sensitive to pass-through. As discussed in the theory, a lower industry-average pass-through exacerbates the market-power distortions and pushes the endogenous carbon price downwards; varying pass-through down to $\tilde{\rho}^M = 0.635$ (see Table 4), we find that the carbon price falls to around \$10 (depending on the precise values of λ, η). By contrast, at its upper estimate $\tilde{\rho}^M = 1.072$ (see Table 4), the qualitative nature of the results flips: because pass-through now exceeds 100%, the industry welcomes a carbon tax and so the endogenous carbon price τ^* exceeds \$30 (Proposition 4). Given our estimates, this case seems relatively unlikely but we also cannot rule it out.

C. WELFARE ANALYSIS AND ENDOGENOUS REGULATION

So far we have focused on the question of how a regulation affects firms' profits. We here extend this analysis in two directions. First, we show how the GLC structure—combined with some additional assumptions—yields results on social welfare. Second, we use these results to endogenize the government's choice on the extent of regulation.

C.1. Welfare analysis. We begin with notation as we now need to be explicit about all firms operating in the industry. Let the set M_i denote the scope of offering of firm $i \in N$ in multi-dimensional GLC (for example, its number of products or markets) so that $\mathbf{x}_i = (x_{im})_{m \in M_i}$, $\mathbf{e}_i = (e_{im})_{m \in M_i}$ are the vectors of its output and factor use. For simplicity, we assume that regulation τ is “complete”, i.e., it applies to the whole offering of all firms (in previous notation, $\Phi \equiv 1$). Unless stated otherwise, we now assume that multidimensional GLC's A1M–A4M apply to all n firms in the industry.

To be able to make statements about welfare, we introduce two additional assumptions:

A5M. Aggregate consumer surplus can be written as $S = V(\mathbf{x}_1, \dots, \mathbf{x}_n) - \sum_{i=1}^n \sum_{m \in M_i} p_{im} x_{im}$, where $V(\cdot)$ is gross consumer utility and consumers are utility-maximizers.

By consumer utility maximization, A5M yields an inverse demand curve for component m of firm i 's offering as $p_{im}(\mathbf{x}_1, \dots, \mathbf{x}_n) = \partial V / \partial x_{im}$. This formulation continues to allow for a wide range of demand-side properties such as substitutes and complements, differentiated products, and so on. It is met, for example, by widely-used linear Bertrand and Cournot models of multiproduct competition (see Section 3.3). As will become clear, we do not require specific functional-form assumptions on gross consumer utility $V(\cdot)$.

A6M. Environmental damages $D(E) \geq 0$ depend on aggregate industry-wide emissions $E = \sum_{i=1}^n \sum_{m \in M_i} e_{im}$ and obey the standard properties $D'(\cdot), D''(\cdot) \geq 0$.

A6M takes the social cost of externality to be global: which firms produce emissions does not matter; only the aggregate is relevant. This is the appropriate setup for climate change and our interest in carbon pricing, and for other global public goods. The setup also allows for there to be no externality, $D(E) \equiv 0$.

Social welfare can then be written in terms of regulation τ as $W(\tau) = S(\tau) + \Pi(\tau) + \tau E(\tau) - D(E(\tau))$, where $\Pi(\tau) \equiv \sum_{i=1}^n \Pi_i(\tau)$ is industry profits and the third term represents government revenue from regulation. Consumer expenditure ($\sum_{i=1}^n \sum_{m \in M_i} p_{im} x_{im}$) is a transfer from consumers to producers while government revenue (τE) is a transfer from producers to government.

Recall that i 's multidimensional rate of cost pass-through is $\rho_i^M(\tau) \equiv \sum_{m \in M_i} \omega_{im}(\tau) \rho_{im}(\tau)$, with weights $\omega_{im}(\tau) \equiv e_{im}(\tau)/e_i(\tau)$, and define an industry-average multidimensional pass-through as $\tilde{\rho}^M(\tau) \equiv \sum_{i=1}^n \frac{e_i(\tau)}{E(\tau)} \rho_i^M(\tau)$.

Our next result characterizes the impact of tighter regulation on consumer surplus and social welfare:

Proposition 3. (a) Suppose that multidimensional GLC's A1M–A3M (not necessarily A4M) hold for each firm $i \in N$ and A5M holds for consumer surplus. Then the first-order change in consumer welfare is given by:

$$\Delta S(\bar{\tau}, \underline{\tau}) \equiv [S(\bar{\tau}) - S(\underline{\tau})] \simeq -(\bar{\tau} - \underline{\tau}) E(\underline{\tau}) \tilde{\rho}^M(\underline{\tau}).$$

(b) Suppose that multidimensional GLC's A1M–A4M hold for each firm $i \in N$ and A5M–A6M hold for consumer surplus and the social cost of the externality. Then the first-order change in social welfare is given by:

$$\Delta W(\bar{\tau}, \underline{\tau}) \equiv [W(\bar{\tau}) - W(\underline{\tau})] \simeq -(\bar{\tau} - \underline{\tau}) \left\{ [1 - \tilde{\rho}^M(\underline{\tau})] E(\underline{\tau}) + [D'(E(\underline{\tau})) - \underline{\tau}] \frac{dE(\tau)}{d\tau} \Big|_{\tau=\underline{\tau}} \right\}.$$

Proof of Proposition 3. For part (a), by A5M, the envelope theorem implies that aggregate consumer surplus satisfies:

$$\frac{dS(\tau)}{d\tau} = - \sum_{i=1}^n \sum_{m \in M_i} x_{im}(\tau) \frac{dp_{im}(\tau)}{d\tau}.$$

Using the definition of i 's firm-level cost pass-through for component m of its offering, $\rho_{im} = [dp_{im}(\tau)/d\tau]/[dk_{im}(\tau)/d\tau]$, and invoking multidimensional GLC's A1M–A3M (but not requiring A4M), we have $dk_{im}(\tau)/d\tau = z_{im}(\tau) = e_{im}(\tau)/x_{im}(\tau)$, again by the envelope theorem, and so this can be rewritten as:

$$\frac{dS(\tau)}{d\tau} = - \sum_{i=1}^n \sum_{m \in M_i} e_{im}(\tau) \rho_{im}(\tau).$$

Finally, we can express this in terms of industry-average pass-through as follows:

$$\frac{dS(\tau)}{d\tau} = - \sum_{i=1}^n e_i(\tau) \sum_{m \in M_i} \frac{e_{im}(\tau)}{e_i(\tau)} \rho_{im}(\tau) = - \sum_{i=1}^n e_i(\tau) \rho_i^M(\tau) = -E(\tau) \sum_{i=1}^n \frac{e_i(\tau)}{E(\tau)} \rho_i^M(\tau) = -\tilde{\rho}^M(\tau) E(\tau).$$

where the second equality uses the definition of i 's multidimensional rate of cost pass-through $\rho_i^M(\tau) \equiv \sum_{m \in M_i} \omega_{im}(\tau) \rho_{im}(\tau)$, with weights $\omega_{im}(\tau) \equiv e_{im}(\tau)/e_i(\tau)$, and the fourth equality uses the definition of industry-average multidimensional pass-through $\tilde{\rho}^M(\tau) \equiv \sum_{i=1}^n \frac{e_i(\tau)}{E(\tau)} \rho_i^M(\tau)$. For part (b), given A5M–A6M, social welfare is $W(\tau) = S(\tau) + \Pi(\tau) + \tau E(\tau) - D(E(\tau))$ and differentiation gives:

$$\frac{dW(\tau)}{d\tau} = \frac{dS(\tau)}{d\tau} + \frac{d\Pi(\tau)}{d\tau} + E(\tau) - [D'(E(\tau)) - \tau] \frac{dE(\tau)}{d\tau}.$$

From Proposition 2(a) for multidimensional GLC, we already know that, given A1M–A4M, i 's firm-level profit impact satisfies $d\Pi_i(\tau)/d\tau = -2[1 - \rho_i^M(\tau)]e_i(\tau)$. Aggregating this across all n firms in the industry then gives:

$$\frac{d\Pi(\tau)}{d\tau} = -2[1 - \tilde{\rho}^M(\tau)]E(\tau),$$

again using the definition of industry-average multidimensional pass-through. Putting together these results yields:

$$\frac{dW(\tau)}{d\tau} = -[1 - \tilde{\rho}^M(\tau)]E(\tau) - [D'(E(\tau)) - \tau]\frac{dE(\tau)}{d\tau}.$$

This establishes the claims using standard first-order approximations.

Proposition 3 establishes first-order welfare impacts based on an industry-average of multi-dimensional pass-through in the spirit of the sufficient-statistics approach (Chetty, 2009; Weyl and Fabinger, 2013). Part (a) is particularly simple: the “static” incidence of tighter regulation is $(\bar{\tau} - \underline{\tau})E(\underline{\tau})$ and the “dynamic” version of consumer welfare declines by the overall proportion that is passed through, $\Delta S \simeq -\tilde{\rho}^M(\underline{\tau})[(\bar{\tau} - \underline{\tau})E(\underline{\tau})]$. This result does not require supply-side linearity (A4M) nor any assumptions about the social cost of the externality (A6M); it relies solely on envelope properties for producers (A1M–A3M) and consumers (A5M).

Part (b) gives an expression that consists of two terms for the impact of tighter regulation on social welfare. The first term captures the combined marginal impact on consumer surplus, industry profits and government revenue. To understand it, consider the case with full pass-through, $\tilde{\rho}^M = 1$: our firm-level Proposition 2(a) implies that industry profits then remain constant while the reduction in consumer welfare is exactly offset by higher government revenue. By contrast, in the case with pass-through below 100%, this term turns negative due to lower industry profits. The second term captures the social value of a reduction in the externality that arises if regulation decreases industry emissions, $dE(\tau)/d\tau < 0$, and the externality is underpriced at the margin, $\tau < D'$.

C.2. Endogenous regulation. Our main GLC analysis assumes that the regulation τ is exogenous. We here extend the analysis to show how regulation—for the case of an emissions tax with “complete regulation” ($\Phi \equiv 1$)—can be endogenized using the GLC structure. This analysis brings together two strands of prior research: (1) an influential literature following Grossman and Helpman (1994) in which firms lobby a government “for sale”; (2) a classic literature following Buchanan (1969) on the optimal design of emissions taxes under imperfect competition in the product market. Taken together, the distortions due to market power and political lobbying mean that this analysis yields a third-best “political equilibrium” emissions tax.⁵³

The basic model setup is as follows. As in Grossman and Helpman (1994), the government cares about social welfare W but also about political contributions by regulated firms. Let $K_i(\tau)$ denote firm i 's political contribution as a function of the emissions price τ . The government's payoff is $U_{\text{gov}}(\tau) = W(\tau) + \lambda \sum_{i=1}^n K_i(\tau)$, where the parameter $\lambda \geq 0$ measures its openness to lobbying, and larger values of λ mean 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 into the government's payoff function, the first-order condition for the “political equilibrium” emissions tax $\tau^*(\lambda)$ is given by:

$$\frac{dU_{\text{gov}}(\tau)}{d\tau} = \frac{dW(\tau)}{d\tau} + \lambda \sum_{i=1}^n \frac{d\Pi_i(\tau)}{d\tau} = 0. \quad (10)$$

We assume that this problem is well-behaved, and focus on the interesting case of an interior solution with $\tau^*(\lambda) > 0$.

The timing of the game is as follows. First, firms choose their contributions K_i . Second, the government sets the emissions price τ . Third, all n firms compete according to multidimensional GLC's A1M–A4M—now taking τ as given, as per A1M. We assume A5M–A6M are met and define $\eta(\tau) \equiv \frac{d \ln E(\tau)}{d \ln \tau} < 0$ as the tax elasticity of industry emissions.

As a benchmark, recall that the standard welfare-maximizing Pigouvian tax under perfect competition—i.e., marginal utility equals price, and price equals marginal cost—is to set the emissions price at the social marginal damage, $\tau^* = D'(E(\tau^*))$; note that this result, like our

⁵³By assumption, the government does not have access to another policy instrument (such as a price control) to directly address market power.

setup, is based on the assumption that consumers are utility-maximizers (Baumol and Oates, 1988).

Proposition 4. Suppose that multidimensional GLC’s A1M–A4M hold for each firm $i \in N$ and A5M–A6M hold for consumer surplus and the social cost of the externality. At an interior solution, the “political equilibrium” emissions tax that maximizes government utility U_{gov} satisfies:

$$\tau^*(\lambda) = \left[\frac{D'(E(\tau))}{1 + \frac{(1+2\lambda)}{-\eta(\tau)} [1 - \tilde{\rho}^M(\tau)]} \right]_{\tau=\tau^*(\lambda)}.$$

Proof of Proposition 4. We pin down the political-equilibrium tax rate using the government’s first-order condition to maximize $U_{\text{gov}}(\tau)$. Given A1M–A4M, it follows from Proposition 2(a) that the impact on industry profits across all n firms satisfies:

$$\frac{d\Pi(\tau)}{d\tau} \equiv \sum_{i=1}^n \frac{d\Pi_i(\tau)}{d\tau} = -2[1 - \tilde{\rho}^M(\tau)]E(\tau),$$

using the definition of industry-average multidimensional pass-through $\tilde{\rho}^M(\tau) \equiv \sum_{i=1}^n \frac{e_i(\tau)}{E(\tau)} \rho_i^M(\tau)$. Given A1M–A6M, we know from Proposition 3(b) that the welfare impact satisfies:

$$\frac{dW(\tau)}{d\tau} = -[1 - \tilde{\rho}^M(\tau)]E(\tau) - [D'(E(\tau)) - \tau] \frac{dE(\tau)}{d\tau}.$$

Putting these parts together in the government’s first-order condition from (10) shows that the political-equilibrium tax rate $\tau^*(\lambda) > 0$ is determined by:

$$\left. \frac{dU_{\text{gov}}(\tau)}{d\tau} \right|_{\tau=\tau^*(\lambda)} = - \left[(1+2\lambda)[1 - \tilde{\rho}^M(\tau)]E(\tau) + [D'(E(\tau)) - \tau] \frac{dE(\tau)}{d\tau} \right]_{\tau=\tau^*(\lambda)} = 0.$$

Using the definition of the emissions elasticity $\eta(\tau) \equiv \frac{d \ln E(\tau)}{d \ln \tau} < 0$ now gives the expression for $\tau^*(\lambda) > 0$ as claimed.

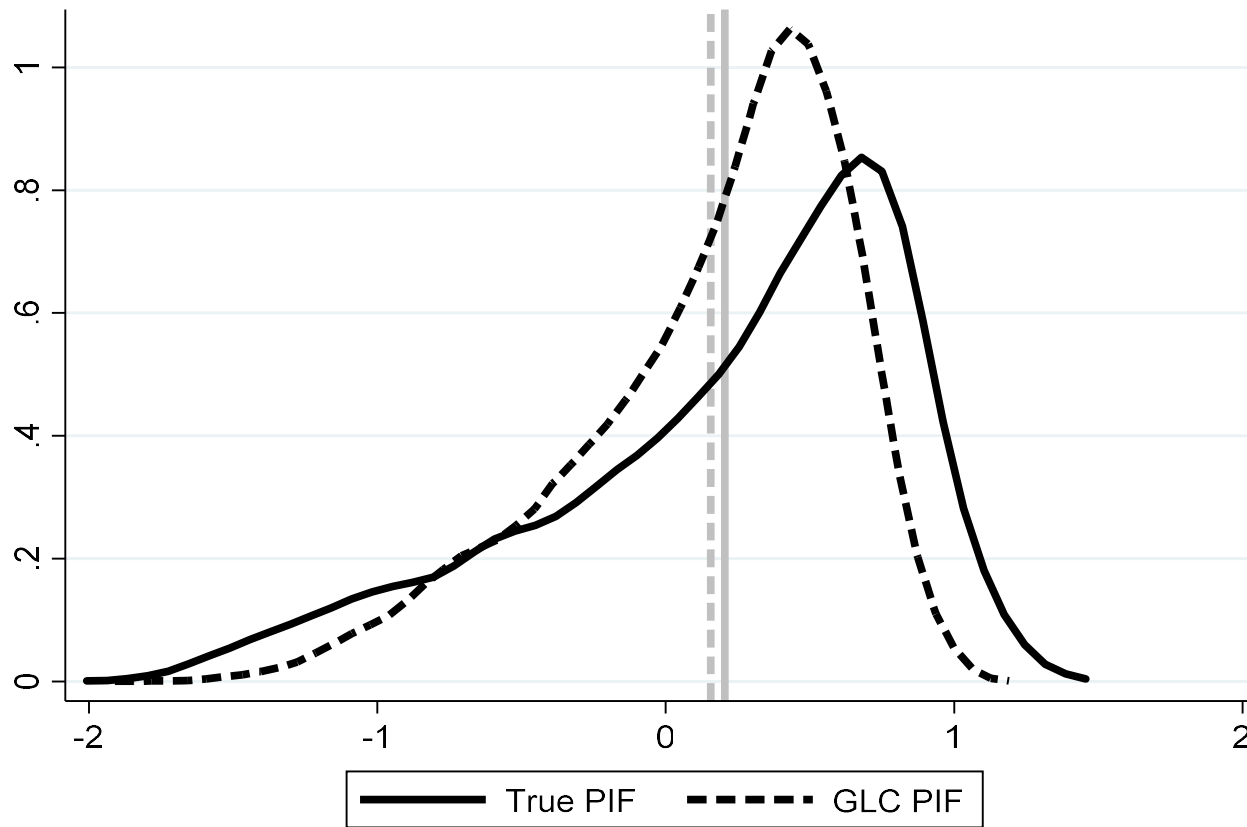
Proposition 4 shows that, for all models belonging to the multidimensional GLC family, the wedge between the equilibrium tax τ^* and the Pigouvian rule is driven by the industry-average multidimensional pass-through rate $\tilde{\rho}^M$. Once again, pass-through signs the result: $\text{sign}(D' - \tau^*) = \text{sign}(1 - \tilde{\rho}^M)$. This generalizes existing literature on the design of second-best emissions taxes under imperfect competition to richer modes of competition and provides a unifying result in terms of pass-through (see Requate, 2006, for an excellent survey of this large literature and its diverse results). It also extends this literature to incorporate Grossman-Helpman political lobbying by regulated firms and thus to third-best emissions pricing.

To understand the result, observe that in the Buchanan problem, industry profits reflect the extent of the market-power distortion while in the Grossman-Helpman problem, profits drive the incentive to make political contributions. We already know from Proposition 2(a) that (first-order) firm-level profit impacts are driven by firm-level multidimensional pass-through—and so the industry-level analog is driven by a weighted average of pass-through across firms.⁵⁴ Intuitively, lower firm-level pass-through ρ_i^M means that firm i contracts output more strongly in response to tighter regulation, creating greater deadweight losses and suffering larger profit losses, thus pushing τ^* downwards—more strongly for large, high-emissions firms (i.e., larger e_i/E). Relatedly, where the government is more open to lobbying (higher λ), and the industry is collectively opposed to the regulation, with $\tilde{\rho}^M < 1$, this pushes the political-equilibrium tax downwards.

⁵⁴Under 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.

In sum, we can therefore also apply Propositions 1 and 2 to calculate the profit impacts of an endogenous regulation, simply by letting $\underline{\tau} = 0$ and $\bar{\tau} = \tau^*(\lambda)$. In the spirit of sufficient statistics, Proposition 4 shows how it is possible to derive an expression for $\tau^*(\lambda)$ that does not directly hinge on knowledge of consumers' utility function $V(\cdot)$ (which, of course, may indirectly matter both for pass-through rates $\tilde{\rho}^M$ and for the emissions elasticity η).

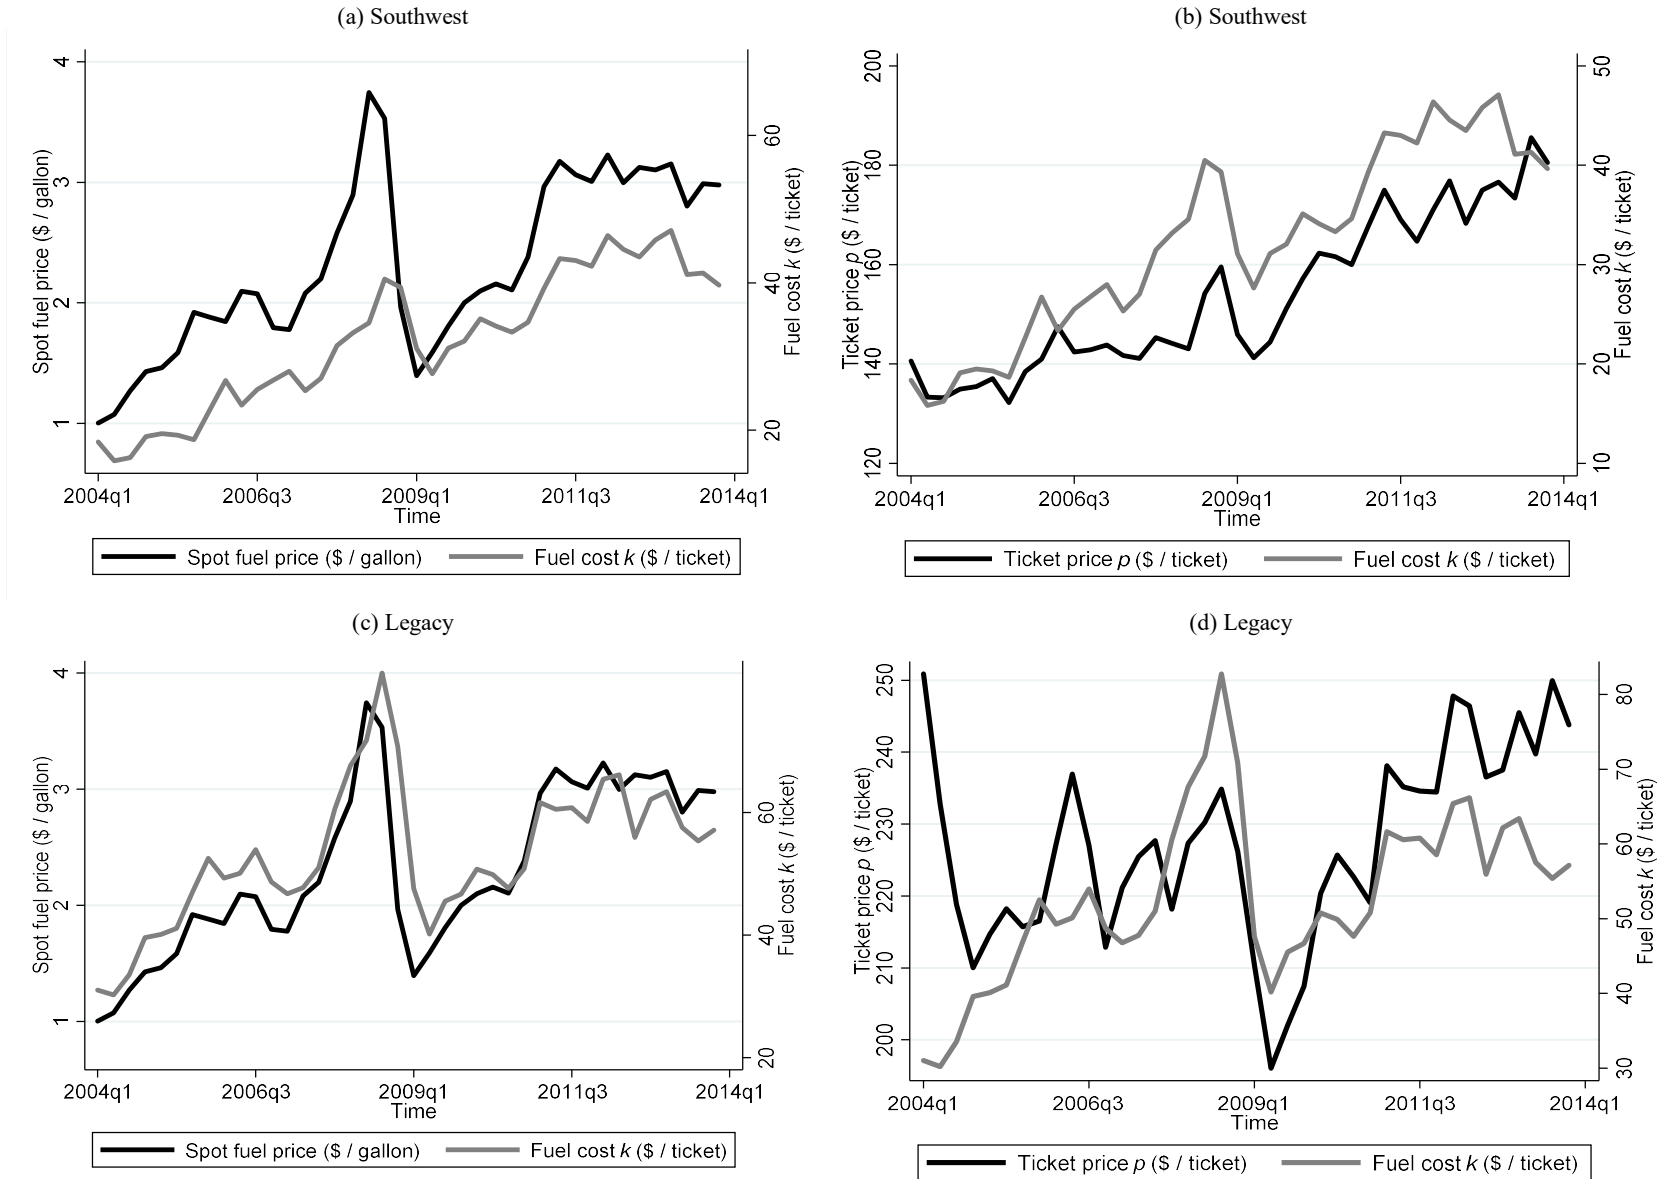
FIGURE 1: DENSITY PLOT OF TRUE AND GLC PROFIT IMPACT FACTOR (PIF)



Notes: Vertical lines in grey indicate the means of the distributions. Kernel densities plotted over all 10,000 draws.

Source: Authors' calculations based on Monte Carlo data described in the main text.

FIGURE 2: AGGREGATE TRENDS IN PRICES AND COSTS



Notes: Panels (a) and (c) show jet fuel spot prices and per-passenger fuel costs; panels (b) and (d) show ticket prices and per-passenger fuel costs. Variables are quarterly averages (unweighted) over all carrier-routes in our sample. Source: Authors' calculations based on data from the US Bureau of Transportation Statistics over the period 2004Q1-2013Q4.

TABLE 1 - PARAMETERS AND RESULTS FROM MONTE CARLO
SIMULATIONS

	(1) Mean	(2) Min	(3) Max
<i>Parameters drawn</i>			
Number of firms n	4.998	2	8
Demand curvature	0.004	-1	1
Cost convexity	0.500	0	1
Market share	0.247	0	0.998
Competitiveness	0.502	0	1
Relative emissions intensity	1.001	0.5	1.5
<i>Results</i>			
True pass-through	0.897	0.316	1.959
GLC profit impact factor	0.205	-1.918	1.367
True profit impact factor	0.156	-1.843	1.123
GLC error	0.049	-0.860	0.523
Number of firm draws	10,000		
Herfindahl Index H	0.317	0.128	0.998

Notes: The table reports parameters and results from Monte Carlo simulations reported in the main text.

Source: Authors' calculations based on Monte Carlo data described in the main text.

TABLE 2 - SUMMARY STATISTICS

	Southwest				Legacy			
	Mean	Standard Deviation	Min	Max	Mean	Standard Deviation	Min	Max
<i>Quarterly av. statistics</i>								
Price (\$)	154.73	40.76	63.07	298.91	227.75	70.89	72.80	599.11
Fuel cost (\$)	32.95	17.86	5.61	129.81	52.98	28.74	2.05	366.63
Distance (miles)	717	466	133	2,298	1,044	629	130	2,724
Emissions (tCO ₂)	0.13	0.06	0.03	0.41	0.20	0.10	0.01	0.71
Emissions cost (\$/tCO ₂)	4.00	1.84	1.03	12.44	6.02	2.89	0.18	21.22
Passengers (000s)	42	39	1	289	30.19	29.21	1.00	244.66
Competitors	2.24	2.40	0	16	2.82	2.29	0.00	16.00
LCC competitors	0.40	0.73	0	4	0.67	0.97	0.00	5.00
Revenue (\$ million)	5.46	4.37	0.12	32.123	5.85	5.84	0.11	53.87
<i>Whole sample statistics</i>								
Revenue in sample (%)	56				43			
Observations	13,199				22,451			
Carrier-routes	416				918			

Notes: Price, fuel cost, emissions and emissions cost are per passenger. Emissions costs are calculated at a carbon price of \$30/tCO₂. Whole sample statistics are aggregated over all N and T. Revenue in sample is the proportion of all US aviation revenue (i.e. all flights on all airlines) in the sample over this period. The legacy carriers are American, Delta, United and US Airways. All averages are unweighted.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE 3 - PASS-THROUGH ESTIMATES

	(1)	(2)	(3)	(4)
Estimation method	OLS	2SLS	2SLS	2SLS
Dependent variable	Price _{ijt}	Price _{ijt}	Price _{ijt}	Price _{ijt}
Sample		Baseline	Balanced panel	Common routes
Fuel cost × Southwest	1.145*** (0.119)	1.418*** (0.121)	1.485*** (0.129)	0.983*** (0.152)
Fuel cost × Legacy	0.560*** (0.082)	0.661*** (0.104)	0.688*** (0.109)	0.518*** (0.140)
Competitors	-2.372*** (0.721)	-2.171*** (0.717)	-2.255*** (0.781)	0.590 (0.996)
LCC competitors	-7.834*** (1.412)	-7.808*** (1.361)	-8.137*** (1.454)	-9.890*** (2.226)
Non-fuel cost index	7.248*** (2.684)	6.895*** (2.580)	7.400*** (2.809)	7.093** (2.993)
GDP growth	-1.374*** (0.398)	-1.341*** (0.388)	-1.620*** (0.436)	0.362 (0.302)
Carrier size	-15.094*** (4.474)	-15.119*** (5.218)	-16.203*** (5.793)	-2.247 (4.502)
Time FE	yes	yes	yes	yes
Carrier-routes FE	yes	yes	yes	yes
Observations	35,650	35,650	24,600	6,138
Clusters	1,334	1,334	615	183
<i>First stage regressions</i>		Dependent variable: Fuel cost × Southwest		
Spot × Southwest		17.258*** (1.065)	17.290*** (1.199)	22.390*** (2.789)
Spot × Distance × Southwest		10.307*** (0.209)	10.486*** (0.219)	9.956*** (0.253)
F-test excluded instruments		1843.69	1545.97	734.71
p-value		(0.000)	(0.000)	(0.000)
		Dependent variable: Fuel cost × Legacy		
Spot × Legacy		25.772*** (3.058)	25.693*** (3.500)	8.412*** (3.058)
Spot × Distance × Legacy		12.090*** (0.820)	12.281*** (0.849)	11.341*** (1.359)
F-test excluded instruments		371.19	374.80	138.67
p-value		(0.000)	(0.000)	(0.000)

Notes: The dependent variable is the price of carrier i , in route j , and quarter t . Standard errors clustered at the carrier-route level are reported in parentheses below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE 4 - PROFIT, CONSUMER SURPLUS AND WELFARE IMPACTS

	(1)	(2)	(3)	(4)	(5)	(6)
	Exogenous $\tau = \$30$			Endogenous $\tau = \$17.68$		
	Southwest	Legacy	All	Southwest	Legacy	All
Pass-through, ρ	1.418 (0.121)	0.661 (0.104)	0.853 (0.109)	1.418 (0.121)	0.661 (0.104)	0.853 (0.109)
Profit impact (in millions \$), $\Delta\Pi$	98.3 [41.4, 155.3]	-233.9 [-377.4, -90.4]	-135.6 [-337.6, 66.4]	58.0 [24.4, 91.5]	-154.3 [-249.0, -59.6]	-79.9 [-198.9, 39.1]
Profit impact (in %), $\Delta\Pi$	1.55% [0.65, 2.45]	-1.60% [-2.59, -0.62]	-0.65% [-1.61, 0.32]	0.91% [0.38, 1.44]	-0.94% [-1.52, -0.36]	-0.38% [-0.95, 0.19]
Consumer surplus (in millions \$), ΔS			-394.8 [-495.8, -293.8]			-232.7 [-292.2, -173.2]
Total welfare (in millions \$), ΔW			49.4 [-51.6, 150.4]			77.3 [17.8, 136.8]

Notes: The first three columns show the profit impact (in absolute terms and as a percentage of revenue) and welfare impacts resulting from an exogenous carbon price of \$30/tCO₂, as implied by the pass-through estimates and Propositions 2(c) and 3. The next three columns give the endogenous carbon price using Proposition 4 and other outcomes using our pass-through estimates. Columns 3 and 6 pass-through rates are an emissions-weighted average of Southwest and legacy results. Standard errors are reported in parentheses below coefficients and 95% confidence intervals in squared brackets.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE B1 - DESCRIPTIVE STATISTICS FOR INDIVIDUAL CARRIERS

	Southwest	American	Delta	United	US Airlines
<i>Quarterly av. statistics</i>					
Price (\$)	154.73	212.54	236.54	231.41	230.85
Fuel cost (\$)	32.95	57.12	50.73	56.31	45.14
Distance (miles)	717	1,102	992	1,128	915
Emissions (tCO ₂)	0.13	0.23	0.18	0.21	0.18
Emissions cost (\$/tCO ₂)	4.00	6.91	5.44	6.23	5.30
Passengers (000s)	42	36	30	28	25.00
Competitors	2.24	2.84	2.4	3.7	2.20
LCC competitors	0.40	0.40	0.73	0.94	0.60
Revenue (\$ million)	5.46	6.76	5.76	5.76	4.66
<i>Whole sample statistics</i>					
Revenue in sample (%)	56%	47%	40%	52%	35%
Observations	13,199	6,110	6,879	5,759	3,703
Carrier-routes	416	198	323	239	158

Notes: Price, fuel cost, emissions and emissions cost are per passenger. Emissions cost are calculated at a carbon price of \$30/tCO₂. Whole sample statistics are aggregated over all N and T. Revenue in sample is the proportion of all US aviation revenue (i.e. all flights on all airlines) in the sample over this period. The legacy carriers are American, Delta, United and US Airways. All averages are unweighted.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE B2 - PASS-THROUGH ESTIMATES FOR INDIVIDUAL CARRIERS

	(1)	(2)
Estimation method	2SLS	2SLS
Dependent variable	Price _{ijt}	Price _{ijt}
Sample	Baseline	Individual airlines
Fuel cost × Southwest	1.418*** (0.121)	1.421*** (0.119)
Fuel cost × Legacy	0.661*** (0.104)	
Fuel cost × American		0.668*** (0.121)
Fuel cost × Delta		1.158*** (0.171)
Fuel cost × United		0.445*** (0.077)
Fuel cost × US		0.767*** (0.094)
Competitors	-2.171*** (0.717)	-2.153*** (0.710)
LCC competitors	-7.808*** (1.361)	-7.678*** (1.312)
Non-fuel cost index	6.895*** (2.580)	10.523*** (2.560)
GDP growth	-1.341*** (0.388)	-1.368*** (0.401)
Carrier size	-15.119*** (5.218)	-14.681*** (5.047)
Time FE	yes	yes
Carrier-routes FE	yes	yes
Observations	35,650	35,650
Clusters	1,334	1,334

Notes: Standard errors clustered at the carrier-route level level are reported in parentheses below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE B3 - PASS-THROUGH ESTIMATES FOR INDIVIDUAL CARRIERS (FIRST STAGE RESULTS)

Sample	(1) Baseline	(2) Individual airlines
	<i>First stage regressions</i>	
Spot × Southwest	17.258*** (1.065)	17.187*** (1.065)
Spot × Distance × Southwest	10.307*** (0.209)	10.304*** (0.209)
F-test excluded instruments	1843.69	635.20
p-value	(0.000)	(0.000)
Spot × Legacy	25.772*** (3.058)	
Spot × Distance × Legacy	12.090*** (0.820)	
F-test excluded instruments	371.19	
p-value	(0.000)	
Spot × American		8.004*** (1.882)
Spot × Distance × American		11.778*** (1.525)
F-test excluded instruments		106.70
p-value		(0.000)
Spot × Delta		6.529*** (1.557)
Spot × Distance × Delta		9.578*** (0.452)
F-test excluded instruments		130.98
p-value		(0.000)
Spot × United		10.438*** (2.088)
Spot × Distance × United		14.557*** (1.836)
F-test excluded instruments		138.50
p-value		(0.000)
Spot × US		5.309*** (0.773)
Spot × Distance × US		9.907*** (0.469)
F-test excluded instruments		374.86
p-value		(0.000)

Notes: Standard errors clustered at the carrier-route level are reported in parentheses below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period

TABLE B4 - ROBUSTNESS OF PASS-THROUGH ESTIMATES

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Estimation method	OLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Dependent variable	Price _{ijt}	Price _{ijt}	Price _{ijt}	Price _{ijt}	Price _{ijt}	Price _{ijt}	Price _{ijt}
Sample		Baseline	Balanced panel	Common routes	Baseline	Balanced panel	Common routes
Fuel cost × Southwest	1.418*** (0.121)	1.450*** (0.124)	1.429*** (0.129)	1.468*** (0.124)	1.364*** (0.125)	1.407*** (0.120)	1.368*** (0.124)
Fuel cost × Legacy	0.661*** (0.104)	0.670*** (0.106)	0.664*** (0.106)	0.679*** (0.105)	0.671*** (0.104)	0.658*** (0.104)	0.665*** (0.102)
Competitors	-2.171*** (0.717)	-2.195*** (0.701)	-2.177*** (0.716)	-2.283*** (0.684)	-2.104*** (0.717)	-2.171*** (0.717)	-2.311*** (0.698)
LCC competitors	-7.808*** (1.361)	-7.819*** (1.357)	-7.841*** (1.378)	-7.757*** (1.348)	-7.770*** (1.354)	-7.693*** (1.384)	-7.564*** (1.345)
Non-fuel cost index	6.895*** (2.580)	6.774*** (2.577)	6.847*** (2.586)	6.902*** (2.555)	7.301*** (2.668)	6.991*** (2.584)	8.081*** (2.447)
Carrier size	-15.119*** (5.218)	-14.884*** (5.244)	-15.226*** (5.166)	-15.115*** (5.190)	-9.270*** (4.977)	-15.120*** (5.211)	-12.697** (5.426)
GDP growth	-1.341*** (0.388)					-1.342*** (0.388)	-1.310*** (0.392)
GDP per capita		1.811 (2.477)					
Population (million)			-0.235 (4.469)				
Gravity demand				12.206 (9.951)			
Network density					-8.398*** (5.063)		
Southwest potential entry						1.861*** (1.515)	
Bankrupt							6.414*** (1.908)
Time FE	yes	yes	yes	yes	yes	yes	yes
Carrier-routes FE	yes	yes	yes	yes	yes	yes	yes
Observations	35,650	35,650	35,584	35,584	35,650	35,650	35,650
Clusters	1,334	1,334	1,326	1,326	1,334	1,334	1,334

Notes: Standard errors clustered at the carrier-route level. Regressions are weighted by total emissions of the carrier-route over the period. Significant levels are denoted: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE B5 - PASS-THROUGH ESTIMATES FROM
A ROUTE-BY-ROUTE ESTIMATION

	(1)	(2)
Estimation method	2SLS	2SLS
Dependent variable	Price _{ijt}	Price _{ijt}
Sample	Southwest	Legacy
Fuel cost	1.267*** (0.046)	0.762*** (0.034)
Competitors	-0.903*** (0.239)	-3.733*** (0.563)
LCC competitors	-0.937*** (0.239)	-3.603*** (0.580)
Non-fuel cost index	42.990*** (2.387)	15.948*** (1.992)
GDP growth	1.349*** (0.099)	2.276*** (0.258)
Carrier size	0.336 (0.965)	-32.321*** (3.191)
Quarterly FE	yes	yes
Observations	9,920	14,680
Carrier-routes	248	367

Notes: Newey-West standard errors that are heteroskedasticity and autocorrelation consistent are reported in parentheses below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE B6 - PASS-THROUGH ESTIMATES FOR
MONOPOLY MARKETS

Estimation method	2SLS
Dependent variable	Price _{ijt}
Sample	Monopoly
Fuel cost	0.600*** (0.228)
Non-fuel cost index	7.463** (3.820)
GDP growth	-0.435 (0.569)
Carrier size	-38.937*** (11.960)
Time FE	yes
Carrier-routes FE	yes
Observations	6,714
Clusters	361
<i>First stage regressions</i>	
Spot	44.758*** (3.987)
Spot × Distance	8.441*** (0.522)
F-test excluded instruments	95.05
p-value	(0.000)

Notes: The dependent variable is the price of carrier *i*, in route *j*, and quarter *t*. Standard errors clustered at the carrier-route level are reported in parentheses below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE B7- DEMAND ESTIMATION FOR DUOPOLY MARKETS

	(1)	(2)	(3)	(4)	(5)	(6)
Estimation method	OLS	2SLS	2SLS	OLS	2SLS	2SLS
Dependent variable	passengers _{ijt}	passengers _{ijt}	passengers _{ijt}	passengers _{kjt}	passengers _{kjt}	passengers _{kjt}
Price _{ijt} (x 10 ⁻³)	-1.687*** (0.220)	-3.948*** (1.120)	17.414 (33.283)	0.748*** (0.218)	2.506 (2.391)	1.189 (1.921)
Price _{kjt} (x 10 ⁻³)	0.382** (0.162)	0.777 (1.280)	1.291 (1.511)	-1.541*** (0.274)	-2.049 (1.638)	-7.125* (4.249)
Price ² _{ijt} (x 10 ⁻³)			-0.046 (0.071)			
Price ² _{kjt} (x 10 ⁻³)						0.011* (0.006)
Instruments		Spot fuel price Spot × Distance	Spot fuel price Spot × Distance		Spot fuel price Spot × Distance	Spot fuel price Spot × Distance
Time FE	yes	yes	yes	yes	yes	yes
Carrier-routes FE	yes	yes	yes	yes	yes	yes
Observations	3,091	3,091	3,091	3,091	3,091	3,091
Clusters	167	167	167	167	167	167

Notes: Newey-West standard errors that are heteroskedasticity and autocorrelation consistent are reported in parentheses below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.

TABLE B8 - DEMAND ESTIMATION FOR TRIOPOLY MARKETS

	(1)	(2)	(3)	(4)	(5)	(6)
Estimation method	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Dependent variable	passengers _{ijt}	passengers _{ijt}	passengers _{kjt}	passengers _{kjt}	passengers _{ijt}	passengers _{ijt}
Price _{ijt} (x 10 ⁻³)	-7.457*** (2.303)	-9.636** (4.608)	1.950 (13.174)	-4.518 (3.360)	-3.567 (3.496)	-3.734 (3.463)
Price _{kjt} (x 10 ⁻³)	3.777*** (1.429)	1.527 (3.819)	2.977 (9.219)	-13.433 (25.399)	2.558* (1.438)	2.933* (1.713)
Price _{ijt} (x 10 ⁻³)	0.118 (0.558)	0.359 (0.680)	-1.725 (1.063)	-1.739 (1.171)	-1.320** (0.661)	-0.738 (2.184)
Price ² _{ijt} (x 10 ⁻³)		0.008 (0.013)				
Price ² _{kjt} (x 10 ⁻³)				0.033 (0.037)		
Price ² _{ijt} (x 10 ⁻³)						-0.001 (0.003)
Time FE	yes	yes	yes	yes	yes	yes
Carrier-routes FE	yes	yes	yes	yes	yes	yes
Observations	275	275	275	275	275	275

Notes: Newey-West standard errors that are heteroskedasticity and autocorrelation consistent are reported in parentheses below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on quarterly data from the US Bureau of Transportation Statistics for the period 2004Q1-2013Q4.