Comparing applied general equilibrium and econometric estimates of the effect of an environmental policy shock*

Jared C. Carbone† Nicholas Rivers‡ Akio Yamazaki§
Hidemichi Yonezawa¶

Version: February 2020

Abstract

We compare the employment effect of the British Columbia carbon tax using two empirical methods: a reduced-form econometric model and counterfactual simulations conducted using an applied general equilibrium (CGE) model. The comparison allows us to test the theory-driven predictions of the CGE model. It also allows us to test the identification strategy of our econometric model. Ex-post, we find statistically and economically significant effects on sectoral employment levels from the carbon tax — with employment falling in the most carbon-intensive sectors and rising in the least carbon-intensive. The CGE model predicts employment responses of very similar sign and magnitude to our econometric estimates. We find no evidence to suggest that our econometric estimates are likely to be undermined by general equilibrium effects in this policy setting. Finally, we explore the use of the econometric estimates to deepen the empirical content of the CGE model.

***JEL* classifications: C68, H23, Q54**
Keywords: climate policy, carbon tax, computable general equilibrium.

---

*We acknowledge support from and cooperation with Christoph Böhringer, Tom Rutherford, Randy Wigle, and Environment Canada in developing the computable general equilibrium model. Rivers acknowledges support from the Social Science and Humanities Research Council Canada Research Chairs program. All authors acknowledge support from Carbon Management Canada. The ideas expressed here are those of the authors who remain solely responsible for errors and omissions.

†Division of Business and Economics, Colorado School of Mines, jcarbone@mines.edu
‡Graduate School of Public and International Affairs and Institute of the Environment, University of Ottawa, nrivers@uottawa.ca
§National Graduate Institute for Policy Studies, a-yamazaki@grips.ac.jp
¶Research Department, Statistics Norway, hidemichi.yonezawa@ssb.no

1 Introduction

In this paper, we treat the implementation of the carbon tax in British Columbia (BC) as a natural experiment and compare the results of econometric estimates of its effects to counterfactual experiments conducted using an applied general equilibrium (CGE) model of the Canadian economy. The comparison allows us to test the theory-driven predictions of the CGE model. It also allows us to test the identification strategy of our econometric model, using the CGE model to indicate whether or not general equilibrium policy responses are likely undermine our attempts at statistical inference in this policy setting.

Researchers frequently consult CGE models to quantify the impacts of prospective environmental policies. For example, the substantial literature on second-best environmental taxation suggests that the design of an economy's pre-existing fiscal system and the interactions of new environmental regulation with it can have substantial effects on the distortionary cost of environment policies (Goulder, Parry, and Burtraw 1997). This research agenda relies heavily on CGE counterfactual simulations to quantify the general equilibrium responses of factor markets that drive the results. Much of the evidence to date on the impacts of carbon mitigation policies — particularly unilateral policies where important general equilibrium effects of policy arise through trade in energy and energy-intensive goods — comes from CGE analysis (Carbone and Rivers 2017). Outside of the domain of environmental policy, CGE models are used widely in public finance, international trade and development applications (Shoven and Whalley 1984).

The main virtue of these models is their theoretical foundation, their ability to capture a full spectrum of channels through which economic theory suggests policy interventions may operate, to generate counterfactual scenarios when policies have no historical analog, and to conduct welfare analysis — the ultimate goal of economic policy analysis in many cases. As tools of empirical analysis, their main shortcoming is their loose connection to empirical evidence based on statistical inference; a number of strong assumptions are required to construct them and key parameter values are typically quantified (or calibrated) using out-of-sample econometric estimates or ad hoc methods. This leads researchers to question the validity of their results.

Testing the validity of these assumptions is challenging, however. The settings in which CGE models are most usefully employed — those in which a policy intervention is expected to generate economy-wide changes to prices and sectoral activity levels — also present special challenges for inference, and the dearth of studies with this objective testifies to this fact. One reason for this is that CGE models are often used to assess prospective policies for which there is little or no historical record to draw on for direct empirical testing. Second, most CGE applications explore the impacts of policy interventions at the national or global scale. In this context, it is difficult to isolate the effects of a policy intervention from other, background events in the economy. Quasi-experimental designs to solve this problem rely on the identification of control units which are similar to the economy under study but which do not receive the policy treatment. Both idiosyncratic and systematic differences in world regions often make this proposition untenable. Furthermore, candidates for control units may be affected indirectly by treatment to the extent a policys effects are transmitted through world markets for factors, goods and services.

When researchers have attempted to validate their models, they have relied on a comparison of the counterfactual changes in outcomes from the CGE model with their analogs in the data measured before and after the economic events of interest — i.e., a first-difference research design (Kehoe 2005; Valenzuela et al. 2007; Beckman, Hertel, and Tyner 2011). The concern with this approach is that any change in the economy which coincides with the policy intervention and also affects the outcomes of interest will be interpreted as the causal effect of the policy, a classic form of omitted variable bias. For example, any world event that affects trade flows and was correlated with the implementation of the North American Free Trade Agreement (the subject of Kehoe 2005) would undermine this research design.

Against this background, the past two decades have seen the rise of experimental and quasi-experimental research designs in the program evaluation literature intended to address exactly the type of omitted variable bias described in the previous example. The BC carbon tax, which we treat as a quasi-experiment in our analysis, plausibly fits this definition. It was introduced in such a way as to make it likely to be exogenous to other events taking place in the BC economy at the time and it was limited to emitters located within BC, allowing for the possibility of comparison of these treated units with emitters outside of BC.

The first part of our analysis exploits this fact to develop a triple-difference estimate of the impacts of the carbon tax on industry employment levels. The preferred model uses industry-level outcomes based on: differences in industry-level carbon intensity, differences in the province in which an industry is located, and differences across time. The model contains industry by time and industry by province fixed effects and a differential treatment effect based on the carbon-intensity of different sectors. Therefore, we identify the effects of the carbon tax from industry-province-time variation in the data. We use publicly-available industry-level economic and carbon emission data from Statistics Canada.

We then compare these estimates to the ex ante predictions of a province-level CGE model of the Canadian economy. Subject to the maintained assumptions of the econometric estimator, this allows us to validate the results of the CGE model using a technique that has the potential to purge the most obvious sources of omitted variable bias present in past validation exercises.[1]

Using the econometric model, we find statistically significant impacts of the carbon tax across sectors that differ in benchmark carbon intensity. The model predicts a pattern in which employment in the most carbon-intensive sectors (e.g., coal production and cement, for example) declines while employment in the least carbon-intensive sectors (e.g., services) increases relative to the average sector. Employment in the most affected sectors declines by 10-15%, while growth sectors expand on the order of 0-5%.

The CGE model produces a similar pattern of changes in employment across sectors in response to the carbon tax. When we regress the CGE predictions on the econometric predictions, the estimated slope coefficient is 0.81-0.85 (depending on weighting scheme adopted[2]). A perfect correspondence between the models would yield a coefficient value of 1. Thus, the CGE model predictions are very close in sign and magnitude to those predicted by the econometric results. We take this as evidence that a 'typical' CGE model can usefully predict the economic response to an environmental policy shock.

There are, however, important reasons to question the validity of quasi-experimental research designs like the one presented here in the evaluation of policies with economy-wide impacts. Their identification relies on finding treatment and control units that isolate the effects of interest. In our application, for example, we must assume that plants located outside BC's borders are unaffected by the carbon tax (the so-called stable unit treatment value assumption (SUTVA)). Yet these plants may be competitors with BC firms in both input and output markets. A loss of competitiveness of carbon-intensive industries in BC due to the new tax could cause output and exports of competitors in the rest of Canada to rise in order to capture displaced demand for BC-produced goods. Similarly, lower demand for BC-based goods may lower wages for workers leading them to seek employment in other provinces. As we have noted, a large literature on the effects of unilateral climate policy — based primarily on CGE counterfactual analysis — emphasizes the importance of trade in energy, energy-intensive goods and basic factors as a source of offsetting changes in economic activity in unregulated regions of the world economy, so-called "carbon leakage" effects, which may undermine the effectiveness of these designs (Carbone and Rivers 2017).

This concern leads to the second part of our analysis, which exploits the fact that CGE models are designed to capture theory's prediction regarding the magnitude and sign of these types of economy-wide interactions. We use these theoretical predictions to test the validity of our econometric research design. Subject to the maintained assumptions of the theory

1. McKitrick 1998 presents a related critique, comparing versions of CGE models in which the key parameters are estimated econometrically, based on CES functions or alternative flexible functional forms.

2. We use benchmark sector output (from the CGE model) for the weights.

underlying the CGE model, we can determine the degree of bias introduced by economy-wide influences in our policy setting — an issue on which quasi-experimental econometric analyses are typically silent.[3]

By using "pseudo-data" generated from the CGE model, we re-estimate our econometric model to test this idea. We find that the econometric estimator predicts a very similar pattern of impacts when the pseudo-data either contains or is purged of general equilibrium effects that could contaminate our control units. This indicates that the SUTVA is expected to hold and thus general equilibrium responses are unlikely to undermine the econometric research design in our policy setting.

With support for the performance of the CGE model and the econometric research design established, we proceed in the final part of the analysis to use the econometric model to calibrate some of the key parameters that sensitivity analysis shows to have an important influence on the results of the CGE counterfactual analysis. In particular, our econometric results suggest that the appropriate trade elasticity in the CGE model is 25 to 50% larger than the value adopted in the existing CGE model.

To our knowledge, ours is the first attempt to compare the performance of a CGE model and a quasi-experimental econometric model. It is also the first attempt to evaluate the performance of a CGE model designed to study the effects of a carbon tax or any other environmental regulation.[4]

The BC carbon tax has a number of features that make it well suited to our purpose. First, it was applied uniformly to the purchase of all fossil fuels for all consumers in the province, which makes it straightforward to simulate in a CGE framework. In contrast, most other climate and environmental policies specify particular technological parameters, which is much more challenging to simulate in a CGE framework. Second, it reached \$30 per tonne CO_2 in 2012 and was first implemented in 2008. Thus, the price signal is likely to be sufficiently strong and the history of the policy sufficiently long that its impacts may be measured in the data. Third, the economy had very little time to anticipate the policy. Only a year passed in between the time the BC government first mentioned that it was contemplating a carbon tax and when it actually came into force. Only five months passed

3. See Chetty 2009, Heckman 2010, Keane 2010 and Kuminoff and Pope 2014 for related discussions of the use of structural models to evaluate treatment effects in program evaluation.

4. There are other attempts to compare ex-ante simulations of the effects of major, environmental regulations with ex-post empirical results. For example, Ellerman et al. (2000) summarizes the ex-ante attempts to estimate Phase-I compliance cost under the U.S. Acid Rain Program and compares them with ex-post estimates derived from an industry survey. Carlson et al. (2000) econometrically estimate cost functions for power plants to generate ex-post compliance costs for the same program and compare them with ex-ante estimates from the U.S. Environmental Protection Agency. In both cases, the ex-ante estimates are not based on CGE analysis and the ex-post empirical measures not derived from a quasi-experimental research design.

in between the time the government made the official announcement about the structure of the policy and the time it was implemented. As a result, it is likely that comparing behavior immediately before and after implementation captures the outcomes of interest. Fourth, the BC carbon tax is revenue neutral. That is, all of the carbon tax revenue is returned to individuals and businesses through reductions of other taxes. The corporate income tax rate and the two lowest personal income tax rates were reduced by 5 percent.[5] The BC carbon tax presents a unique opportunity to evaluate a specific set of general equilibrium responses generated by CGE models that are central to the theory of second-best environmental taxation. Fifth, British Columbia introduced the carbon tax unilaterally. The specter of a loss of competitiveness in pollution-intensive, trade-exposed domestic industries and of carbon leakage (i.e., offsetting increases in emissions outside of the regulated jurisdiction) has played an important role in preventing many countries from adopting emission controls. CGE models with descriptions of the system of interregional trade form the principal piece of evidence on the magnitude of the competitiveness and leakage effects (Carbone and Rivers 2017). Thus, our experiments have the potential to shed light on this set of issues as well.

Our paper builds on existing literature that explores the economic impacts of the carbon tax in British Columbia (see Murray and Rivers (2015) for a summary). Beck et al. (2015) and Beck, Rivers, and Yonezawa (2016) implement a computable general equilibrium model to estimate the effect of the carbon tax on household income. Neither of these studies, however, report the estimated impacts of the policy on sector employment or activity levels. Yamazaki (2017) uses an econometric approach to estimate the impact of the carbon tax on sector level employment, while Yip (2018) uses a similar approach to estimate the impact of the tax on individual-level labour force attachment and wages. The current study compares the impacts on employment estimated using a CGE approach to those estimated with a reduced-form econometric approach. Our CGE model builds on Beck et al. (2015), while the econometric model builds on Yamazaki (2017). We aggregate the data used for each model in order to make the models comparable and focus the article on comparing outcomes from the two approaches.

We focus on comparing sectoral changes in employment across our two models. There are a number of reasons to choose this indicator. First, examining the impact of a carbon tax on sectoral activity is of interest to policymakers, since this is closely related to the concepts of competitiveness and emissions leakage, both of which are important considerations for the development of cost-effective and politically acceptable carbon policy (Carbone and Rivers

5. In addition, to protect low-income households, the government gives them a lump-sum credit. In 2012-2013, \$1.4 billion was credited back to individuals and businesses. In fact, tax credits exceeded the tax revenue by \$260 million. This excess is estimated to decline to only \$20 million in 2013-2014.

2017). Second, impacts on employment are themselves central to the political acceptability of climate policy and there is only limited empirical evidence on the effect of a carbon tax on employment (Hafstead and Williams III 2018; Yamazaki 2017). Third, employment is relatively straightforward to measure without error, so it is a useful indicator of the sector-level impact of the carbon tax.

The rest of the paper proceeds as follows. In the following section, we describe both the econometric and computable general equilibrium models that are used, in addition to the sources of data. In section 3, we present results from each model describing the economic response to the introduction of a carbon tax, and also formally compare the results of the two models. In section 4, we focus on uncertainties in the two models, and explore whether it is possible to use the models jointly to narrow the uncertainties. Finally, section 5 concludes.

2 Models

2.1 Econometric model

We econometrically estimate changes in employment in response to the introduction of the carbon tax in British Columbia using a triple-difference framework combined with a treatment intensity variable (akin to a triple differences estimator). We compare differences in these variables (i) at different levels of the carbon tax (and before and after the carbon tax was introduced), (ii) in British Columbia and other provinces, and (iii) in sectors with high and low carbon intensity. Given the assumptions described in the prior section, the triple-difference framework allows us to isolate the causal effect of the carbon tax on sectoral performance in British Columbia. Our approach builds on prior work by Yamazaki 2017 and Rivers and Schaufele 2015.

Formally, our approach for estimating the impact of the carbon tax on employment is to estimate:

$$\ln L_{irt} = \beta_1(EI_{ir} \times \tau_{rt}) + \beta_2\tau_{rt} + \lambda^1_{ir} + \lambda^2_{it} + \alpha X_{rt} + \epsilon_{irt} \tag{1}$$

where L_{irt} is employment of industry i in region/province r at time t, EI_{ir} is the emissions intensity of industry i in region r, measured in tonnes of greenhouse gases per dollar of output,[6] τ_{rt} is the value of the carbon tax in province r in year t, λ are fixed effects, and ϵ_{irt} is an idiosyncratic error term.

Successful identification of the β parameters of interest is contingent on the fixed effects absorbing potentially confounding variables. λ^1_{ir} is an industry-province fixed effect that absorbs the average employment by industry in each province. λ^2_{it} is an industry-time fixed

6. We explain the measurement of emissions intensity in a later section.

effect that absorbs any common shocks by industry and year, for example as a result of changes in commodity prices or national policy.

Threats to identification of β_2 come in the form of disturbances that are correlated with τ_{rt} and not absorbed by the fixed effects. For example, if British Columbia (or other provinces) implemented other province-wide policies concurrently with the carbon tax, or if there were shocks to labour supply in British Columbia that varied over time relative to other provinces, estimates of β_2 would be biased. In order to guard against non-parallel trends, we estimate a version of the model that includes linear province time trends, thus identifying the difference from pre-existing trends in employment across provinces. In addition, we estimate a version of the model that controls for a number of observable variables that vary at the province-year level and that are potentially correlated with (but not caused by) the carbon tax. These are given by X_{rt}.

Threats to identification of β_1 come in the form of disturbances that are correlated with $\mathrm{EI}_{ir} \times \tau_{rt}$. Identification of β_1 therefore requires much weaker assumptions than identification of β_2: to generate bias in β_1, omitted variables must be correlated not only with the carbon tax, but with the linear combination of the tax and sector emission intensity. The assumption that the impact of the carbon tax varies linearly with sector emission intensity is not derived from a theoretical model, but serves an important purpose in identifying the impact of the tax.[7]

2.2 Computable general equilibrium model

We use a static multi-sector, multi-region computable general equilibrium (CGE) model of the Canadian economy to simulate changes in sectoral employment in response to the introduction of the carbon tax in British Columbia. The model is a "standard" implementation of an energy-focused computable general equilibrium model. The model has previously been used in several other applications for assessment of climate change policy in Canada.[8] This

7. An alternative approach would be to estimate a separate difference-in-difference specification for each industry, and drop the imposed assumption that the impact of the carbon tax varies linearly with emission intensity. However, this approach would come at a cost, since in such an approach omitted variables correlated with the carbon tax would generate bias in sector-level point estimates. We view the linearity assumption as much milder than the potential omitted variables problem, and thus use the triple-difference approach to estimate the impact of the carbon tax.

It is also worth noting that the presence of omitted variables that are correlated with τ_{rt} but *uncorrelated* with EI_{ir} are not sufficient to introduce bias in estimates of β_1. Because the β_2 term is included in the model as well, β_1 is identified using the differential effect of the carbon tax across sectors of different emission intensities.

8. For example, see Böhringer et al. 2015 for an application to burden-sharing, Böhringer, Rivers, and Yonezawa 2016 for an application to fiscal federalism, and Beck et al. 2015 for an application related to the incidence of carbon taxes.

section includes a non-technical overview of the model. A more formal model description is provided in an online appendix to this article.

The model captures characteristics of provincial (regional) production and consumption patterns through detailed input-output tables and links provinces via bilateral trade flows. Each province is explicitly represented as a region, except Prince Edward Island and the Territories, which are combined into one region. The representation of the rest of the world is reduced to import and export flows to Canadian provinces which are assumed to be price takers in international markets. To accommodate analysis of energy and climate policies the model incorporates rich detail in energy use and greenhouse gas emissions related to the combustion of fossil fuels.

The model features a representative agent in each province that receives income from three primary factors: labour, capital, and fossil-fuel resources.[9] There are three fossil resources specific to respective sectors, namely, coal, crude oil and gas. Fossil-fuel resources are specific to fossil fuel production sectors in each province. Labour is treated as perfectly mobile between sectors within a region, but not mobile between regions. We assume that half of the capital stock is mobile between sectors and provinces, while the other half is specific to each sector in each province. The model incorporates details of direct and indirect taxes which are received by the provincial or federal governments in order to finance public services.

The construction of a multi-sectoral CGE (such as ours) involves a choice regarding which sectors present in the underlying calibration data to represent explicitly and which to combine into more aggregated sectors in the model to maintain computational tractability. Because our policy application focuses on carbon mitigation, we leave carbon-intensive sectors disaggregated. Industrial carbon emissions are tightly linked to energy use. Thus, we also explicitly represent primary fossil energy sectors and the electricity sector. The energy goods identified in the model include coal, gas, crude oil, refined oil products and electricity. This disaggregation is essential in order to distinguish energy goods by carbon intensity and the degree of substitutability. In addition, the model features major energy-intensive industries which are potentially those most affected by emission reduction policies.

We describe in some detail the structure of the model with a focus on sectoral production and trade in the sub-sections below. We focus on these aspects of the model since they are the key model structures that affect the response to carbon policy, which we compare with the econometric evidence. Further details on other aspects of the model are available at Böhringer et al. 2015.

9. Labour supply is endogenous, following the specification in Ballard 2000. Uncompensated (compensated) elasticity of labour supply is 0.05 (0.3). Land use associated with agricultural production and forestry is not explicitly accounted for, but instead treated as part of the specific capital stock of the relevant sector.

Production

Production of commodities in each region and sector pair (Y_{jr}) is captured by multi-level constant elasticity of substitution (CES) production functions combining capital, labour, energy and materials.[10] Production occurs via constant-returns-to-scale firms operating in perfectly competitive markets.[11] At the top level of the production function, a composite of non-energy intermediate material demands (M_{jr}) is combined with an aggregate of energy, capital, and labour (KLE_{jr}) subject to a constant elasticity of substitution (σ^M):[12]

$$Y_{jr} = \left(\theta_{jr}^{M}(M_{jr})^{\frac{\sigma^M-1}{\sigma^M}} + \theta_{jr}^{KLE}(KLE_{jr})^{\frac{\sigma^M-1}{\sigma^M}}\right)^{\frac{\sigma^M}{\sigma^M-1}} \tag{2}$$

where θ_{jr}^{M} is the value share of intermediate inputs (M) in the production of Y. θ is similarly defined in the equations below.

The intermediate good composite is a fixed (Leontief) composite of the M individual intermediate inputs, each of which is an Armington composite of imports and domestic production (as described below):

$$M_{jr} = \min(\theta_{jr}^{M_1} A_{1r}^{j}, \ldots, \theta_{jr}^{M_M} A_{Mr}^{j}) \tag{3}$$

where A_{kr}^{j} is Armington good k in region r used as an input to sector j ($k = j$). Other inputs include energy, capital, and labour. In the main model specification, capital and labour are aggregated into a CES composite (V_{jr}), which is then combined with an energy composite (E_{jr}):

10. From these production functions, we derive the CES cost functions describing the price-dependent use of capital, labour, energy and materials, which are included in the zero profit equations shown in the more formal model description in the online appendix.

11. There is a growing literature focused on market power and its effect on the pass-through of energy-related costs. See Ganapati, Shapiro, and Walker, Forthcoming for example. Intuitively, greater market power leads to less pass-through of costs to consumers, resulting in a less elastic output response in an industry. Some industries — such as the gasoline industry in the U.S. — have been found to have pass-through rates significantly below 1, which could dampen the employment response to a carbon tax relative to an otherwise-equivalent competitive industry. In principle, it would be possible to formulate a CGE model which exhibits market power and could be calibrated with estimates of pass-through rates. However, we have chosen not to consider this complication in our comparison. This is because a meaningful comparison would require estimating an extension of our econometric model which contains carbon-policy interactions with measures of industry market power. Measures of market power at the industry-province level for Canada do not exist to our knowledge. National measures (such as those reported in Crépeau and Duhamel 2008) would not contain enough variation to identify these parameters in our framework.

12. We drop subscripts on σ^M and other elasticities of substitution in the production function to reduce clutter, although these are differentiated by sector as described below.

$$KLE_{jr} = \left(\theta^E_{jr}(E_{jr})^{\frac{\sigma^E-1}{\sigma^E}} + \theta^V_{jr}(V_{jr})^{\frac{\sigma^E-1}{\sigma^E}}\right)^{\frac{\sigma^E}{\sigma^E-1}} \tag{4}$$

where

$$V_{jr} = \left(\theta^K_{jr}(K_{jr})^{\frac{\sigma^L-1}{\sigma^L}} + \theta^L_{jr}(L_{jr})^{\frac{\sigma^L-1}{\sigma^L}}\right)^{\frac{\sigma^L}{\sigma^L-1}} \tag{5}$$

Values for the elasticities of substitution between capital and labour, as well as between value-added (V_{jr}) and energy differ by sector and are drawn from the econometric work of Okagawa and Ban 2008.[13]

The energy composite is a nested CES aggregate of electricity, gas, refined petroleum products (oil), and coal. Specifically, the aggregate energy input is defined as a CES function of electricity and the composite of coal, oil and gas; and the composite, coal, oil and gas is a CES function of coal and a CES aggregate of oil and gas:

$$E_{jr} = \left(\theta^{ele}_{jr}(ele_{jr})^{\frac{\sigma^{ELE}-1}{\sigma^{ELE}}} + (1-\theta^{ele}_{jr})(CGO_{jr})^{\frac{\sigma^{ELE}-1}{\sigma^{ELE}}}\right)^{\frac{\sigma^{ELE}}{\sigma^{ELE}-1}} \tag{6}$$

where:

$$CGO_{jr} = \left(\theta^{col}_{jr}(col_{jr})^{\frac{\sigma^{COA}-1}{\sigma^{COA}}} + (1-\theta^{col}_{jr})(GO_{jr})^{\frac{\sigma^{COA}-1}{\sigma^{COA}}}\right)^{\frac{\sigma^{COA}}{\sigma^{COA}-1}} \tag{7}$$

and

$$GO_{jr} = \left(\theta^{gas}_{jr}(gas_{jr})^{\frac{\sigma^{OIL}-1}{\sigma^{OIL}}} + (1-\theta^{gas}_{jr})(oil_{jr})^{\frac{\sigma^{GO}-1}{\sigma^{OIL}}}\right)^{\frac{\sigma^{OIL}}{\sigma^{OIL}-1}}. \tag{8}$$

Elasticities of substitution for energy goods are $\sigma^{ELE} = 0.25$, $\sigma^{COA} = 0.5$, and $\sigma^{OIL} = 0.75$. These elasticities take on similar values in other CGE models (see e.g., Paltsev et

13. In the simulations below, we test the impact of an alternative nesting structure in which labour is combined with a CES aggregate of capital and energy, also based on econometric estimates from Okagawa and Ban 2008. We also test the impact of replacing the econometrically-estimated elasticities from Okagawa and Ban 2008 with those from Dissou, Karnizova, and Sun 2015. This results in a total of four alternative production function specifications (two different nesting structures and two different sets of econometric estimates).

al. 2005). They are also consistent with substantial empirical evidence available in this regard (e.g., (Stern 2012)).

In the production of fossil fuels (coal, crude oil and natural gas), the production function is similar to that described above, except the capital-labour-energy-materials aggregate is combined with a fossil fuel specific resource at the top level. The elasticity of substitution between this sector-specific resource and the other inputs is calibrated to reflect empirical evidence on fossil fuel supply elasticities as described in Rutherford 2002.

In all of the simulations we consider, we take technology as exogenous. That is, firms can move along isoquants in response to changes in (relative) prices, but isoquants are fixed. This assumption effectively rules out innovation as a response to changes in emission prices.

Trade

Bilateral trade between provinces as well as between each province and the rest of the world is specified following the Armington (1969) approach, which distinguishes domestic and foreign goods by origin. For example, the machinery produced in British Columbia is an imperfect substitute for the machinery produced in other Canadian provinces and the rest of the world. Since we assume that the aggregation of the domestic supply (from Canadian provinces) and the import from the rest of the world is modelled as a nested CES function, the relative prices of domestic supply and import determine each demand quantity with the given substitution elasticities.[14] Specifically, at the top nest, we combine the domestically produced good (or the composite of the goods produced in the Canadian provinces) and the imported good from the rest of the world.

$$A_{jr} = \left(\theta^A_{jr} (H_{jr})^{\frac{\sigma^{DM}-1}{\sigma^{DM}}} + (1-\theta^A_{jr})(M_{jr})^{\frac{\sigma^{DM}-1}{\sigma^{DM}}} \right)^{\frac{\sigma^{DM}}{\sigma^{DM}-1}},$$

where H_{jr} is good j produced domestically (within Canada) that is consumed in province r, and M_{jr} is imports of good j to province r from the rest of the world. σ^{DM} is the elasticity of substitution between domestic goods and imports. At the second nest, we combine the good produced in the own province (D^r_{jr}) and goods produced in other provinces s (D^s_{jr}), and σ^{PP} is the elasticity of substitution between own-province good and goods produced in other provinces.[15]

14. This will be clear by deriving cost function from the CES production function and then deriving import demand function by applying the envelope theorem to the cost function (see the model description in the online appendix).

15. To be more precise, in our simulation model, we also have the trade margin for each bilateral trade among provinces, which can be interpreted as iceberg trade costs, and they are combined to the bilateral

$$H_{jr} = \left(\theta_{jr}^{H_r} (D_{jr}^r)^{\frac{\sigma^{PP}-1}{\sigma^{PP}}} + \sum_{s \neq r} \theta_{jr}^{H_s} (D_{jr}^s)^{\frac{\sigma^{PP}-1}{\sigma^{PP}}} \right)^{\frac{\sigma^{PP}}{\sigma^{PP}-1}} .$$

Regarding the export, we assume the constant elasticity of transformation (CET) function to split the production into the domestic supply (to the provinces in Canada) and to the rest of the world. Thus, the quantities of domestic supply and international export are determined by the relative prices of domestic supply and export.[16]

$$Y_{jr} = \left(\theta_{jr}^{D} (D_{jr})^{\frac{\eta-1}{\eta}} + (1-\theta_{jr}^{D})(X_{jr})^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} ,$$

where D_{jr} is production of good j in province r that is sold on the domestic market (both own province and other provinces) and X_{jr} is the volume sold to the rest of the world. η is the elasticity of transformation between domestic goods and exports.

All Canadian provinces are assumed to be price takers in the world market. There is an imposed balance of payment constraint between Canada and the rest of world aggregate. To implement this constraint, we fix the current account surplus exogenously at the benchmark level.

The elasticities η, σ^{DM}, and σ^{PP} determine the changes in exports and imports as a function of changing relative prices of domestic and foreign goods. In the benchmark model calibration, we set $\eta = \sigma^{DM} = 4$ and set $\sigma^{PP} = 2\sigma^{DM}$. We do now know of a source that empirically estimates Armington elasticities for Canadian provinces, and so these parameters are chosen to roughly reflect empirical literature from other regions (Donnelly et al. 2004). We set σ^{PP} at twice the level of σ^{DM} to reflect the 'border effect' which suggests that goods move more easily over provincial and state borders than over national borders (Anderson and Van Wincoop 2003). A similar approach is taken in Caron, Rausch, and Winchester 2015.

Simulation

Using the CGE model, we conduct an analysis of the impact of the introduction of the carbon tax in British Columbia. The tax was phased in over five years starting in 2008, and reached a value of \$30/t CO_2 in 2012. This tax value was applied on all combustion greenhouse

trade at the fixed proportion (or Leontief function).

16. In a similar way to the derivation of the import demand function, export supply function can be derived from the CET transformation function.

gas emissions in the province, while non-combustion emissions were not covered by the tax (in total, the tax covered about 75% of total emissions). We ignore transitional dynamics associated with the tax's introduction, and focus on its 2012 level (BC's government has stated that the tax will be held at this level until at least 2018). To analyze the effect of the tax, we construct two model simulations: one in which the tax is applied to all fossil fuel combustion in British Columbia, and one in which no tax is applied. We infer the effect of the tax in the CGE model by comparing these two counterfactual simulations.

British Columbia's carbon tax was introduced as part of a broader environmental fiscal reform in which revenues from the tax were used in part for lump-sum transfers to low-income households, and in part for reducing rates of pre-existing personal and corporate income taxes in the province. We replicate the revenue recycling scheme that was implemented by the province in our analysis, as described in detail in Beck et al. 2015. One feature of the implementation of the tax is that government carbon tax revenues were slightly lower than associated tax rebates, such that the deficit position of the provincial government increased slightly in conjunction with the tax. We replicate this change in our simulation of the tax.

2.3 Data sources and reconciliation

In order to make meaningful comparisons between the two modeling approaches, we need to ensure that data used for both models is comparable. This section explains the sources of data used for each model, and the steps taken to ensure comparability.

The CGE model adopts the calibration approach to parameterization in which cost and income share parameters are drawn from the benchmark social accounting matrix and free elasticity parameters are drawn from estimates in the published literature (as described above). Statistics Canada input-output and final demand data provides the foundation for the social accounts matrix (Statistics Canada 2006a, 2006b). This data is not available at a high sectoral resolution at the provincial level, and so we make use of disaggregated national data as well as proprietary data from Environment Canada to further disaggregate certain energy-intensive sectors, as well as to obtain data on sectoral greenhouse gas emissions. We use economic data from 2007 (prior to introduction of the carbon tax) as the benchmark data source.

Table 1 summarizes key features of the resulting data set underlying the CGE model. There are 17 sectors in each region of the model, including three primary energy sectors (natural gas, crude oil, and coal), two secondary energy sectors (refined petroleum products and electricity), five manufacturing sectors (pulp and paper, primary metals, chemicals, cement, and other), as well as several other sectors that are generally less energy- and emissions-intensive. Note that the emissions intensity includes both the direct emissions (which are associated with fossil fuel combustion in the sector) as well as the indirect emissions (which are associated with fossil fuel combustion in the production of non-fuel goods and services that are used as intermediate inputs in the sector).[17] We calculate indirect emissions based on the benchmark input-output matrix, using the method described in Rutherford 2010. We only include emissions covered by the BC carbon tax in this calculation (i.e., due to fossil fuel combustion in BC), since this is the basis for estimating the impact of the tax.[18]

The largest sectors in the economy produce few emissions, notably the service, government, construction, other manufacturing, and trade sectors. The sectors that produce the largest amount of emissions per unit output are the coal mining and cement, followed by the crude oil production, primary metal manufacturing, and transport sectors. Manufactured goods and energy goods are highly traded. Figure 1 visually summarizes key sectoral data, and shows the sectors potentially most 'exposed' to carbon prices: those with high emissions intensity as well as a large export share (which likely limits ability to pass through the carbon cost to final consumers). These sectors include coal mining, cement, and the crude oil extraction sectors, and to a lesser extent the heavy manufacturing sectors (chemicals, pulp and paper), the transport sector, and the natural gas production sector. Each of these represents a small share of the total BC economy.

The econometric approach is based on data on industry-level employment collected by Statistics Canada's Survey on Employment, Payrolls, and Hours (SEPH). We use annual data by province and NAICS industry.[19] The Statistics Canada data reports the total number of employees in each industry, province, and year. While the CGE model uses data on 17 sectors (indexed by j in the discussion above) per region, the econometric model uses data on about 80 industries (indexed by i). We map these disaggregate industries to the more aggregate CGE sectors, as shown in online appendix Table A.1.

Our econometric analysis also employs data on the emissions intensity (and trade intensity) of each industry. Unfortunately, Statistics Canada does not maintain data on the emissions intensity of industries disaggregated at a provincial level. In a similar analysis, Yamazaki 2017 uses the national emissions intensity as a proxy for the emissions intensity for each industry and region. We build on that approach here, but also make use of the provincial sector-level emissions intensity from the CGE model database. This is important

17. For example, in the metal sector, the combustion of natural gas generates direct emissions, whereas the usage of electricity results in indirect emissions because electricity consumes natural gas. Note that while electricity generation in BC is mainly hydro, a small portion of electricity is generated by natural gas.

18. In other words, the indirect emissions that are associated with the imported intermediate inputs are not included because the BC carbon tax is not imposed on those emissions.

19. Statistics Canada. Table 14-10-0202-01 Employment by industry, annual (formerly CANSIM series 281-0024).

since in some cases there are significant differences between national emissions intensity in a given sector and the corresponding provincial measure (e.g., in electricity generation, where some provinces produce virtually no emissions and others produce significant quantities per unit of electricity generated). In particular, we construct the emissions intensity for industry i in sector j in province r based on national Statistics Canada industry-level data ($\tilde{EI}_i$) and provincial-sector emissions intensity from the National Accounts ($\bar{EI}_{jr}$) as follows:

$$EI_{ijr} = \tilde{EI}_i \times \frac{GDP_{jr} \times \bar{EI}_{jr}}{\sum_{i \in j} GDP_{ijr} \times \tilde{EI}_i}$$

This measure preserves the intra-industry variation in emissions intensity that is observed in the national data, but ensures that aggregate sector emissions intensity matches the provincial data. Using this measure ensures that the econometric results are directly comparable to the results generated by the CGE model. Thus, we replace EI_{ir} with EI_{ijr} in Eq.(1).

3 Results

3.1 Econometric model results

Table 2 shows the results corresponding to the estimation of Eq.(1). The coefficient on τ_{rt} ($\hat{\beta}_2$ in Eq.(1)) is estimated based on overall differences in employment in British Columbia after the carbon tax was implemented with employment prior to the implementation of the tax. The coefficient on $(EI_{ijr} \times \tau_{rt})$ ($\hat{\beta}_1$ in Eq.(1)) captures heterogeneity in the effect of the tax across industries due to differences in emissions intensity. Different columns in Table 2 correspond to the inclusion of different fixed effects and control variables in the estimation of Eq.(1).

In the first column of the table, we include fixed effects for industry by time and industry by province. These fixed effects respectively control for unobserved heterogeneity and events that vary systematically by industry over time (for example, changes in trade policy, changes in technology, or changes in commodity prices), as well as controlling for unobserved factors that cause some industries to be larger in some provinces than others (for example, different resource bases, policy, or preferences). Identification of the main coefficients in the model is contingent on there being no other unobserved heterogeneity that is correlated with the treatment variables in the model. The results in the first column suggest that sectors with higher carbon intensity are more adversely affected by the imposition of the carbon tax (the coefficient on $(EI_{ijr} \times \tau_{rt})$ is negative). However, the coefficient on τ_{rt} is positive, suggesting that for sectors with low enough emission intensity, the net effect on employment is positive.[20]

In subsequent columns of the table, we successively include a more complete set of fixed effects and province-time varying control variables to control for unobserved time-varying shocks that could bias the results. In column (2), we include several time-varying provincial controls including a measure of the demographic distribution of the population (i.e., the population in each of seven exhaustive age groups),[21] and interactions between province dummy variables and the world oil price, the US unemployment rate, and the volume of world trade. The variables are included to capture potentially heterogeneous responses across provinces to shocks to world trade and energy markets, such as due to the great recession. This could be important, because, for example, some provinces are oil importers while others are oil exporters, so that a change in the world price of oil would be expected to impact provinces differently.

In column (3), we also include province-specific linear and quadratic time trends to capture time varying unobservables that differ across provinces. Our preferred specification, is given in column (3), and includes industry-province fixed effects and industry-time fixed effects, province-time controls and province time trends as described above. We select this as our preferred specification because it includes a full set of fixed effects as well as an extensive set of control variables to capture heterogeneity in provinces over time.

Finally, in column (4), we include province-time fixed effects. Province-time fixed effects absorb potentially unobserved confounders that vary by province and year, such as other provincial policies. These fixed effects are collinear with the introduction of the carbon tax (which also varies by province and year) and so it is not possible to identify β_2 in this specification. Instead, we use it as a robustness check for β_1. It is encouraging that the estimate of β_1 changes very little from column (1)-(3) to column (4), suggesting that this coefficient is not biased by unobserved covariates that vary over time within provinces.

In column (3), which we focus on, the model suggests that sectors with higher emissions intensity are more likely to experience reductions in employment as a result of the carbon tax than sectors with low emissions intensity. The coefficient on $(EI_{ijr} \times \tau_{rt})$ suggests that for a sector with emission intensity of 1 tonne CO_2 per \$1,000, a \$1 increase in the carbon tax is associated with a 0.52 percent reduction in employment compared to a sector with no emissions.[22] The coefficient is estimated precisely, and is similar in magnitude to Yamazaki

20. Industries with emission intensity below (above) $-\beta_2/\beta_1$ experience positive (negative) employment impacts from the carbon tax.

21. By including this covariate, we are implicitly making the assumption that the carbon tax does not influence demographic outcomes. We do not have evidence on this point, but removing the variable has little impact on coefficient estimates.

22. The emissions intensity measure is described above. Briefly, we sum direct and indirect emissions covered by the carbon tax originating in BC and divide by gross output.

2017 who uses a similar approach and data. The coefficient on τ_{rt} suggests that for a sector with no emissions, employment increases by 0.1 percent when the carbon tax increases by \$1. This coefficient is not statistically significantly different from zero at standard levels. This is perhaps unsurprising, as the overall effect of the carbon tax on employment is likely to be small, and the number of observations where the tax takes on different values over time and by province is also small. These two coefficients can be used to generate predicted values for the impact of the carbon tax on employment. For example, employment in the Primary Metal manufacturing sector (which has emissions intensity of 0.619 tCO_2/\$1,000) falls by about 6.3% as a result of the carbon tax at \$30/t CO_2.[23]

Our predicted values from the econometric model are shown as blue triangles in the left-hand panel of Figure 2. Results are ordered by their emission intensity (right-hand panel). The econometric model predicts a reduction in employment in the emissions-intensive sectors of the economy, and a growth in employment in the less emissions-intensive sectors. Importantly, the largest sectors of the economy — services, wholesale and retail trade, other manufacturing, government, and construction — all have low emissions intensity (the middle panel shows the size of each sector). Conversely, the emission-intensive sectors such as the cement, coal, and primary metals sectors represent a much smaller fraction of the overall economy.

3.2 Computable general equilibrium model results

In the left-hand panel of Figure 2 depicts the predicted impact of the carbon tax on sector employment levels in British Columbia from our CGE model using red circles. The sectors predicted by the model to experience the most significant impacts in terms of employment losses are the coal and cement sectors, which the model suggests reduce employment by 17% and 12%, respectively. In addition, the model suggests employment losses of 7 to 10% in the petroleum refining, primary metal manufacturing, and transport sectors. As shown in the right panel of Figure 2, these sectors have high emission intensity relative to other sectors in BC. In contrast, the model suggests increases in employment in the agriculture, construction, other manufacturing, services, and retail and wholesale trade sectors, all of which are low greenhouse gas intensity sectors (impacts are ordered by sector greenhouse gas intensity, which is shown in the right panel of Figure 2).

23. We calculate predicted values using $\Delta \hat{L}_j = \exp(30 \times (\hat{\beta}_1 \times EI_j + \hat{\beta}_2)) - 1$.

3.3 Comparison of results

Figure 3 plots the sectoral predictions from the econometric model and the CGE model against each other. In the figure, the dashed black line has an intercept of zero and a slope of 1, such that all points would fall along this line if the two models were identical. There is good visual concordance of the results between the two approaches. Sectors for which the CGE model predicts a large negative impact — in particular the coal mining and cement sectors — are also predicted to experience a large negative impact based on the econometric results. Sectors for which the CGE model predicts will experience a benefit as a results of the carbon tax — in particular the low emission sectors such as services, wholesale and retail trade, and other manufacturing - are also predicted to experience a benefit based on the econometric results.

Table 3 describes a formal comparison of the results. In this table, we show results produced by regressing the results from the CGE model on the point estimates of predictions generated using the econometric model. We conduct both an unweighted regression and a weighted regression. Predicted values from each regression are shown in Figure 3 as dashed red and blue lines, and slope coefficients from each regression are given in Table 3. Figure 3 highlights the similarity of results derived from the regression model and CGE model. In particular, sectors that are more predicted to be more significantly impacted by the tax in the econometric model are generally also predicted to be more impacted by the tax in the CGE model. Table 3 shows that the slope coefficient from a regression of the CGE predicted values on the econometric predicted values is around 0.81 to 0.85, depending on whether coefficients are weighted, again indicating a close concordance between the two approaches.

It is also worth noting that while the slope of the regression models is close to one, the intercept is somewhat below zero. This relates to the β_2 coefficient in Eq.(1). Unlike β_1, the variable relating to this coefficient varies at the province-time level, such that unobserved province-time heterogeneity could bias the coefficient estimate.

In each case, the R^2 value from the regressions suggest that the predicted impacts from the econometric model are highly predictive of the impacts generated from the CGE model. We also formally test the hypothesis that the slope of the regression of econometric predictions of CGE simulations is 1 and do not reject ($p = 0.26$), and test the joint hypothesis that the slope is 1 and the intercept is 0 and fail to reject ($p = 0.20$). Overall, the comparison reveals that the two sets of models predict very similar outcomes.

4 Robustness testing

Our comparison of the econometric and CGE models suggests that the two models predict similar patterns of sector-level employment response to a carbon tax. In this section we extend our analysis by using the models jointly to explore the validity of the assumptions in each.

Our econometric strategy relies on three key assumptions. First, we assume that industries in control provinces are unaffected by the introduction of the BC carbon tax. If industries are connected through interprovincial trade or factor markets, this may not be the case. Offsetting impacts in control industries (i.e. increases in employment) in response to the tax could also cause our model to overestimate the treatment effect due to contamination of control units.

Second, the baseline specification of the econometric model also fails to control for differences in trade exposure across industries. All else equal, theory predicts that more trade-exposed industries should be more heavily impacted by the carbon tax. Leaving the role of trade exposure as the genesis of contamination effects aside, it is possible that — if a few trade-exposed industries are disproportionately impacted by the tax — they could skew the point estimates in our baseline model upward.

Third, our baseline econometric model also assumes that the impact of the carbon tax is linear in the carbon intensity of a sector of the economy. To the extent that this aspect of the model is misspecified, our point estimates could be biased.

Our CGE model also contains a large number of assumptions. Rather than trying to conduct an exhaustive sensitivity analysis of the model, our strategy for validation is to focus on some of the key parametric assumptions that are closely connected to carbon-policy outcomes of interest in our experiments.

Our first test focuses on the trade elasticity values assumed in the model. Past CGE studies of unilateral carbon policy consistently find that policy responses are sensitive to these parameter values because they determine the degree to which carbon-intensive production can re-locate across the global economy in response to increased production costs in the regulated region.

Our second test focuses on the substitution elasticities and production nesting structures that characterize demand for carbon-based energy in the model — another central determinant of the impacts of carbon policy. Mulptile nesting structures and accompanying sets of estimated elasticity values have been proposed in the literature, so we explore how this range of possibilities influences our core findings.

4.1 Testing for contamination effects

The first test we perform is aimed at evaluating the tenability of the assumption that control industries are unaffected by the carbon tax (the so-called stable unit treatment value assumption (SUTVA)). This is normally a maintained assumption in the type of empirical analysis we present here, and it is not straightforward to check for the validity of the assumption. In our case, we are able to use the CGE model to probe the likely validity of the SUTVA (with the test obviously being conditional on the assumptions in the CGE model).

To explore the importance of SUTVA violations, we construct two pseudo-data based on the output from the benchmark and counterfactual equilibria produced by our CGE model. The first data contains all of the general-equilibrium responses produced by the counterfactual experiment and allows us to capture the ways in which they contaminate outcomes for control provinces in our triple-difference model. We purge the second data of these influences for control provinces to establish a benchmark against which to judge the magnitude of these contamination effects.

The first data uses the benchmark-equilibrium outcomes from the CGE model to populate the pre-treatment periods and the counterfactual-equilibrium outcomes for the post-treatment periods. This procedure applies to outcomes for both treatment and control provinces.

In the second data, we populate the treatment province (BC) outcomes using the benchmark-equilibrium outcomes for the pre-treatment periods and the counterfactual-equilibrium outcomes for the post-treatment periods. We populate the control province outcomes using the benchmark-equilibrium outcomes for both the pre-treatment periods and the post-treatment periods. Thus, we assume there is no change in quantity or price variables in control regions post-treatment.

Because our CGE model is static (i.e. it captures a snapshot of the economy at a single point in time), there are no time trends to account for in the construction of the pseudo-data. Therefore, the counterfactual response to the carbon tax is the only source of time variation they contain. Furthermore, the only factor that could invalidate the triple-difference estimator in this setting is contamination of control-province outcomes due to GE responses.

Table 4 displays the results of our regressions using the pseudo-data. The format of the tables mirrors the presentation of our baseline econometric results, and the empirical strategy and identification strategy are identical. Aside from the contents of the data used, there is a difference between the baseline regressions and the pseudo-regressions presented here. As previously noted, the level of aggregation is considerably higher in the CGE model than in the data used to estimate the econometric model. In this exercise, we use the more aggregate data generated by the CGE model.

Table 4 contains three columns. In the first column, we estimate Eq.(1) using the pseudo-data generated by the CGE model. The interpretation of the coefficients is the same as in our baseline econometric specification. Due to the similarity of the results from the econometric and CGE models previously described, we find coefficients of the same sign and very close to the same magnitude as when the original data is used. Specifically, sectors with higher carbon intensity are expected to cut employment more significantly when the carbon tax is imposed. By combining the real data and pseudo-data, we can formally test the equality of the coefficients from estimating the model on the two different data sets, and are not able to reject equality of the two sets of coefficients.[24]

In the second column, we estimate Eq.(1) using the pseudo-data generated from the CGE model, but purging any spillover effects of the carbon tax on counterfactual units – sectors and regions outside of BC. The coefficients change only very slightly in response to this change, suggesting that the SUTVA is roughly maintained in the CGE model. To the extent that the CGE model is a valid representation of the real world, we would expect a similar spillover effect, and thus would expect the SUTVA to also hold. In the third column, we include a variable that measures the trade intensity of each sector (proportion of output that is exported) and interact this variable with the carbon tax rate. The coefficient is small and statistically insignificant, while other coefficients are unchanged. This suggests that differences in trade intensity are not a key factor explaining heterogeneous employment impacts of the carbon tax, and that the inclusion of this variable in our main econometric model is not warranted (a point we further explore below).

4.2 Controlling for trade-exposure

One might expect the offsetting increases in employment in control industries would be a problem primarily in goods that are heavily traded across provinces — so that when output is reduced in BC under the carbon tax, suppliers from other provinces can step in to satisfy demand. Here we explore an alternative regression in which we add interaction terms between our main carbon-tax treatment and the trade intensity of different sectors in the model. The resulting specification becomes:

$$\ln L_{ijrt} = \beta_1(EI_{ijr} \times \tau_{rt}) + \beta_2\tau_{rt} + \beta_3(TI_{ijr} \times \tau_{rt}) + \lambda^1_{ijr} + \lambda^2_{ijt} + \epsilon_{ijrt} \tag{9}$$

where TI_{ijr} is defined as the share of output that is exported in industry i of sector j in province r in 2007.

24. An F-test fails to reject equality of the two interaction coefficients p=0.78.

With the addition of the trade-intensity interaction, the original two treatment terms (β_1 and β_2) now measure the impact of the carbon tax on sector employment conditional on trade exposure. If the contamination of control industries runs through trade — or if failing to control for trade exposure otherwise results in biased estimates of these coefficients — then examining the results of this regression may give an indication of the severity of these problems.

We report the regression results in column (2) of Table 5. The coefficients on the original variables (β_1 and β_2 in Eq.(1)) are not markedly changed from the original estimates, which are presented in column (1), suggesting that the absence of trade exposure did not significantly bias the results. The coefficient on the trade intensity variable itself is not significantly different from zero, and is of unexpected sign (i.e., we would expect that a higher trade intensity would be associated with reduced employment in the presence of the tax; this is not what we found).

We obtain some support for the notion that sectoral trade intensity is not a strong determinant of the impact of the tax from similar estimation using the CGE pseudo-data. In particular, a regression including an interaction between the carbon tax and sector trade intensity (alongside other variables) does not yield an intuitively signed or precisely estimated coefficient, and other coefficients magnitudes do not change as a result of the inclusion of the trade-intensity variable (column (3) of Table 4).

Based on these results, it appears that an econometric specification that omits sector trade intensity is likely appropriate.[25] Interestingly, this is the same finding as reported in Martin et al. 2014. And importantly, it suggests that existing policies, which compensate industries for carbon pricing based partly on trade intensity (such as the EU and California), are likely not targeting compensation optimally.

4.3 Evaluating functional forms

Our baseline econometric model assumes that the impacts of the carbon tax treatment are linearly related to the carbon intensity of a sector's production technology. However, clearly there is scope for alternative functional forms, and thus we also estimate an alternative, non-parametric version of the econometric model with sector-specific treatment terms as shown in Eq.(10).

$$\ln L_{ijrt} = \sum_j \beta_j(\text{Sector}_j \times \tau_{rt}) + \lambda^1_{ir} + \lambda^2_{it} + \alpha X_{rt} + \epsilon_{ijrt} \tag{10}$$

25. While in general we do not find that trade exposure is an important determinant of outcomes, in the case of crude oil - which is extremely trade-exposed in our setting - the CGE model suggests that it may be important.

where Sector_j is a sector indicator variable. λ^1_{ir} is an industry-province fixed effect while λ^2_{it} is an industry-time fixed effect. X_{rt} is the same control variable included in our baseline specification of the model.

It is important to note that relaxing the linearity assumption comes at the expense of imposing much stronger identification assumptions. Whereas the original econometric specification uses a triple-difference type specification, the specification in Eq.(10) uses a difference-in-difference specification, since it does not rely on comparing outcomes across sectors with different emission intensity (instead estimating a different treatment effect for each sector). As a result, it is vulnerable to threats to identification that stem from the impact of events in BC that coincide with the introduction of the carbon tax. For this reason, we prefer to use the main specification in our analysis.

The results of these regressions are imprecise, but mostly show similar effects to our baseline results. Figure 4 summarizes the results from the regressions on employment. While the standard errors are large due to larger data demands of the non-parametric specification, they also show a clear trend in the point estimates that mirrors our baseline results — with the most carbon-intensive industries seeing the largest declines in employment on a scale that is approximately linear with carbon intensity. Figure 5 illustrates the concordance of the results between the CGE and the non-parametric econometric approach. Results are mostly similar, with the exception of two sectors (CHM - chemical manufacturing and MIN - mining).

4.4 Parameterizing the CGE model

Eq.(2) through Eq.(8) describe the functional forms adopted for production and trade in the CGE model, which determine the sectoral responses to the carbon tax examined here. Choice of how to parametrize these functional forms can be key to determining the sectoral outputs. To determine the sensitivity of the results to changes in parametrization, we conduct two sensitivity analyses. First, we test the sensitivity of the results to changes in trade elasticities: σ^{D_i} and σ^{A_i}. We conduct one run where these elasticities are doubled from their initial levels, and another in which they are halved.

Second, we test the sensitivity of the results to changes in the production function. The base parameterization of the model adopts a production function where capital and labour are combined in one nest, and this aggregate is then combined with an energy aggregate. Elasticities of substitution are from Okagawa and Ban 2008. Okagawa and Ban 2008 also estimate elasticities for a version in which capital and energy are in a nest, and this nest is combined with labour, and we test this alternative functional form with their estimated elasticities. Dissou, Karnizova, and Sun 2015 conduct similar estimation using Canadian data, for both of these nesting structures. In total, this gives four alternative sets of (sector-specific) production functions, and we run our model using each.

Figure 6 shows the sensitivity of the model results to differences in the trade elasticities, and Figure 7 shows the sensitivity of the model results to differences in the production function. The CGE model is somewhat sensitive to changes in the trade elasticities. In particular, for sectors in which a significant reduction in employment is predicted, the trade elasticity has an important effect. Higher trade elasticities reduce the ability of the sector to pass through costs of the carbon tax, and exacerbate losses in emissions-intensive sectors. Differences are especially noteworthy in the cement sector, and to a lesser degree in the oil refining, primary metal manufacturing, and transport sectors. In contrast, the structure of the production function and the elasticities in the production function play very little role in determining the effect of the carbon tax on sectoral activity (at least, within the bounds of the combinations examined here).

Given the sensitivity of the CGE model to changes in the specification of trade elasticities, as well as the limited empirical basis for specifying these elasticities (recall that there is no data available on Canadian provincial trade elasticities), we attempt to use the results of the econometric model to verify and improve the specification of trade elasticities in the CGE model. In particular, we conduct simulations in the CGE model using alternative trade elasticity settings, and compare the results of these simulations to the econometric results, using the methods described earlier. We alter the trade elasticities in the model by multiplying the elasticities η, σ^{DM}, and σ^{PP} by a common multiplier (set at 1.0 in the original model specification). We show the results of this experiment in Figure 8. The original setting, described earlier, results in a R-squared value of 0.74 from a regression of the CGE results on the econometric results, and a slope coefficient of 0.83 from the same regression. It is clear that model fit improves somewhat when we increase the trade elasticities in the model. In particular, the slope coefficient is one when trade elasticities are set at 1.5 times their original value, and the overall model fit is highest when trade elasticities are set at about 1.3 times their original value.

5 Conclusions

Here we have compared ex-ante estimates of the effects of the BC carbon tax based on a detailed CGE model of the Canadian economy with ex-post estimates derived from a quasi-experimental econometric model. This allows us to test the theory and calibration underlying the CGE model as well as the potential for general equilibrium effects to undermine our econometric research design in this policy setting. Overall, we find a strong degree of correspondence in the sign and relative magnitude of the sectoral impacts predicted by the two models and support from theory for our econometric model.

While the experiments described here represent only a single case study, they suggest that the CGE model is a useful tool for making ex-ante predictions about the economic effects of environmental policies and for making welfare calculations. They also demonstrate how these models can play complementary roles in the evaluation of large-scale environmental policies — as a framework for model validation and in the use of statistical inference to deepen the empirical content of CGE analysis. An important role for future research in this area should be to develop empirical tests of CGE models in other policy contexts and to systematically explore threats to validity in research designs used in the program-evaluation literature using CGE models.

References

Anderson, James E., and Eric Van Wincoop. 2003. "Gravity with gravitas: A solution to the border puzzle." *The American Economic Review* 93 (1): 170–192.

Ballard, Charles. 2000. "How many hours are in a simulated day? The effects of time endowment on the results of tax-policy simulation models." *Unpublished, Michigan State University.*

Beck, Marisa, Nicholas Rivers, and Hidemichi Yonezawa. 2016. "A rural myth? Sources and implications of the perceived unfairness of carbon taxes in rural communities." *Ecological Economics* 124:124 –134.

Beck, Marisa, Nicholas Rivers, Randall Wigle, and Hidemichi Yonezawa. 2015. "Carbon tax and revenue recycling: Impacts on households in British Columbia." *Resource and Energy Economics* 41:40–69.

Beckman, Jayson, Thomas Hertel, and Wallace Tyner. 2011. "Validating energy-oriented CGE models." *Energy Economics* 33 (5): 799–806.

Böhringer, Christoph, Nicholas Rivers, and Hidemichi Yonezawa. 2016. "Vertical fiscal externalities and the environment." *Journal of Environmental Economics and Management* 77:51–74.

Böhringer, Christoph, Nicholas Rivers, Thomas Rutherford, and Randall Wigle. 2015. "Sharing the burden for climate change mitigation in the Canadian federation." *Canadian Journal of Economics/Revue canadienne d'économique* 48 (4): 1350–1380.

Carbone, Jared C., and Nicholas Rivers. 2017. "Climate policy and competitiveness effects: Policy guidance and quantitative evidence." *Review of Environmental Economics and Policy* 11 (1): 24–42.

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer. 2000. "Sulfur dioxide control by electric utilities: What are the gains from trade?" *Journal of Political Economy* 108 (6): 1292–1326.

Caron, Justin, Sebastian Rausch, and Niven Winchester. 2015. "Leakage from sub-national climate policy: the case of California's cap-and-trade program." *The Energy Journal:* 167–190.

Chetty, Raj. 2009. "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods." *Annu. Rev. Econ.* 1 (1): 451–488.

Crépeau, Stéphane, and Marc Duhamel. 2008. *Competition intensity in Canada: A critique of recent OECD findings.* Technical report 2008-09. Industry Canada.

Dissou, Yazid, Lilia Karnizova, and Qian Sun. 2015. "Industry-level econometric estimates of energy-capital-labor substitution with a nested CES production function." *Atlantic Economic Journal* 43 (1): 107–121.

Donnelly, W.A., K. Johnson, M.E. Tsigas, and D.L. Ingersoll. 2004. *Revised Armington elasticities of substitution for the USITC model and the concordance for constructing consistent set for the GTAP model.* Technical report 20001-A. United States International Trade Commission.

Ellerman, A. Denny, Paul L. Joskow, Richard Schmalensee, Elizabeth M Bailey, and Juan-Pablo Montero. 2000. *Markets for clean air: The U.S. Acid Rain Program.* Cambridge University Press.

Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker. Forthcoming. "Energy cost pass-through in US manufacturing: Estimates and implications for carbon taxes." *American Economic Journal: Applied Economics.*

Goulder, Lawrence H., Ian W.H. Parry, and Dallas Burtraw. 1997. "Revenue-raising versus other approaches to environmental protection: The critical significance of preexisting tax distortions." *RAND Journal of Economics* 28:708–731.

Hafstead, Marc A.C., and Roberton C. Williams III. 2018. "Unemployment and environmental regulation in general equilibrium." *Journal of Public Economics* 160:50–65.

Heckman, James J. 2010. "Building bridges between structural and program evaluation approaches to evaluating policy." *Journal of Economic Literature* 48 (2): 356–398.

Keane, Michael P. 2010. "Structural vs. atheoretic approaches to econometrics." *Journal of Econometrics* 156 (1): 3–20.

Kehoe, Timothy J. 2005. "An Evaluation of the Performance of Applied General Equilibrium Models on the Impact of NAFTA." In *Frontiers of Applied General Equilibrium Modelling.* Cambridge, UK: Cambridge University Press.

Kuminoff, Nicolai V., and Jaren C. Pope. 2014. "Do capitalization effects for public goods reveal the public's willingness to pay?" *International Economic Review* 55 (4): 1227–1250.

Martin, Ralf, Mirabelle Muûls, Laure B. de Preux, and Ulrich J. Wagner. 2014. "On the empirical content of carbon leakage criteria in the EU Emissions Trading Scheme." *Ecological Economics* 105:78–88.

McKitrick, Ross R. 1998. "The econometric critique of computable general equilibrium modeling: the role of functional forms." *Economic Modelling* 15:543–573.

Murray, Brian, and Nicholas Rivers. 2015. "British Columbia's revenue-neutral carbon tax: A review of the latest grand experiment in environmental policy." *Energy Policy* 86:674 –683.

Okagawa, Azusa, and Kanemi Ban. 2008. "Estimation of substitution elasticities for CGE models." *Discussion Papers in Economics and Business* 16.

Paltsev, Sergey, John M. Reilly, Henry D. Jacoby, Richard S. Eckaus, James R. McFarland, Marcus C. Sarofim, Malcolm O. Asadoorian, and Mustafa H.M. Babiker. 2005. *The MIT emissions prediction and policy analysis (EPPA) model: version 4.* Technical report. MIT Joint Program on the Science and Policy of Global Change.

Rivers, Nicholas, and Brandon Schaufele. 2015. "The effect of carbon taxes on agricultural trade." *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie* 63 (2): 235–257.

Rutherford, Thomas F. 2010. "Climate-linked tariffs: Practical issues." Unpublished manuscript prepared for: Thinking ahead on international trade, 2nd Conference on Climate Change, Trade, and Competitiveness: Issues for the WTO, Graduate Institute of International and Development Studies.

———. 2002. "Lecture notes on constant elasticity functions." Unpublished manuscript, Department of Economics, University of Colorado.

Shoven, John B., and John Whalley. 1984. "Applied general equilibrium models of taxation and international trade: An introduction and survey." *Journal of Economic Literature* 22 (3): 1007–1051.

Statistics Canada. 2006a. *Final demand by commodity, S-Level Aggregation, Table 381-0012.*

———. 2006b. *Inputs and Output by industry and commodity, S-Level Aggregation, Table 381-0012.*

Stern, David I. 2012. "Interfuel substitution: A meta-analysis." *Journal of Economic Surveys* 26 (2): 307–331.

Valenzuela, Ernesto, Thomas W. Hertel, Roman Keeney, and Jeffrey J. Reimer. 2007. "Assessing Global Computable General Equilibrium Model Validity Using Agricultural Price Volatility." *American Journal of Agricultural Economics* 89, no. 2 (): 383–397.

Yamazaki, Akio. 2017. "Jobs and climate policy: Evidence from British Columbia's revenue-neutral carbon tax." *Journal of Environmental Economics and Management* 83:197 – 216.

Yip, Chi Man. 2018. "On the labor market consequences of environmental taxes." *Journal of Environmental Economics and Management* 89:136 –152.

6 Figures

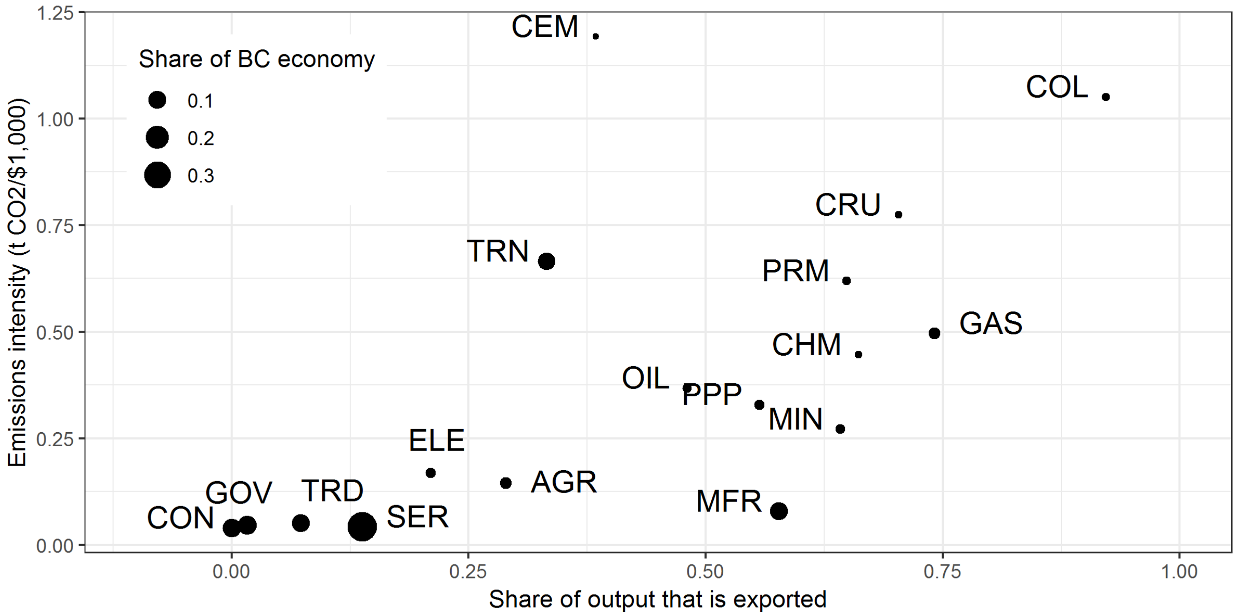

Figure 1: Benchmark CGE data. The emissions intensity is the direct and indirect emissions that are subject to the BC carbon tax (i.e., originating from fossil fuel combustion in British Columbia) divided by sector gross output. Area of circles corresponds to the size of each sector relative to the entire BC economy.

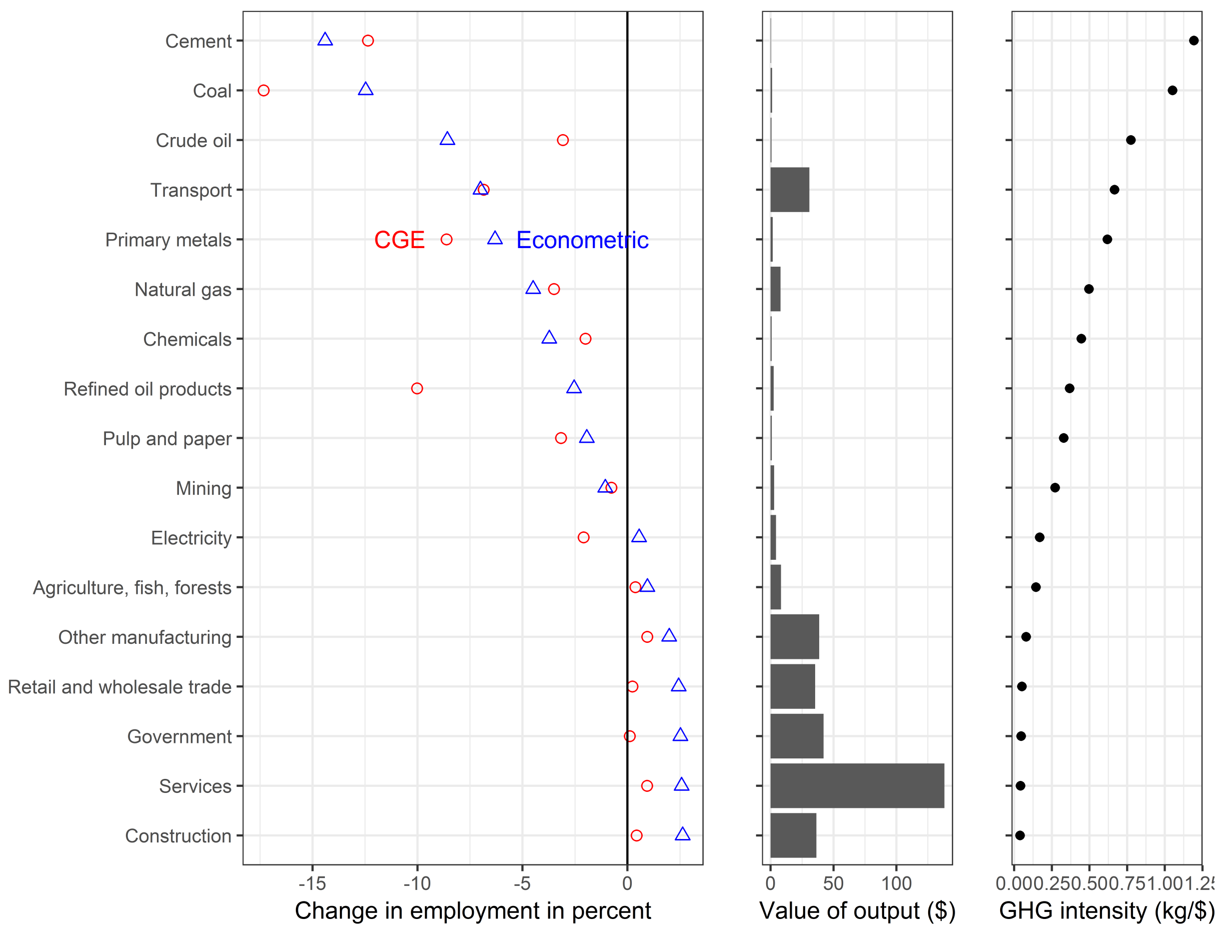

Figure 2: Left-hand panel: Point estimates of predicted change in employment by sector due to a \$30/t CO_2 tax based on econometric estimates (blue triangle) and CGE model (red circle). Middle panel: Total sector output in benchmark. Right-hand panel is GHG emission intensity.

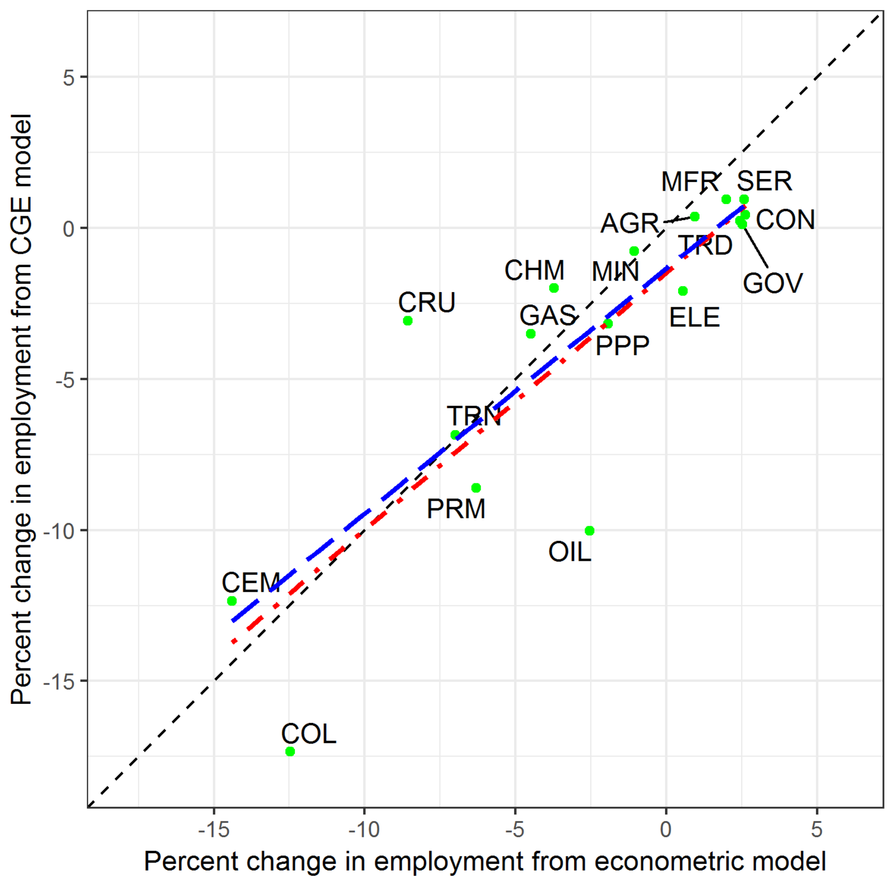

Figure 3: Comparison of econometric and CGE estimates of change in employment by sector associated with unilateral adoption of a \$30/t CO_2 tax in British Columbia. Dashed black line has intercept of 0 and slope of 1; points would fall along this line if the two models were identical. The red dash-dot line is the regression line from an unweighted regression of predicted values from the CGE model on predicted values from the econometric model. The blue long-dash line is the regression line from a weighted regression.

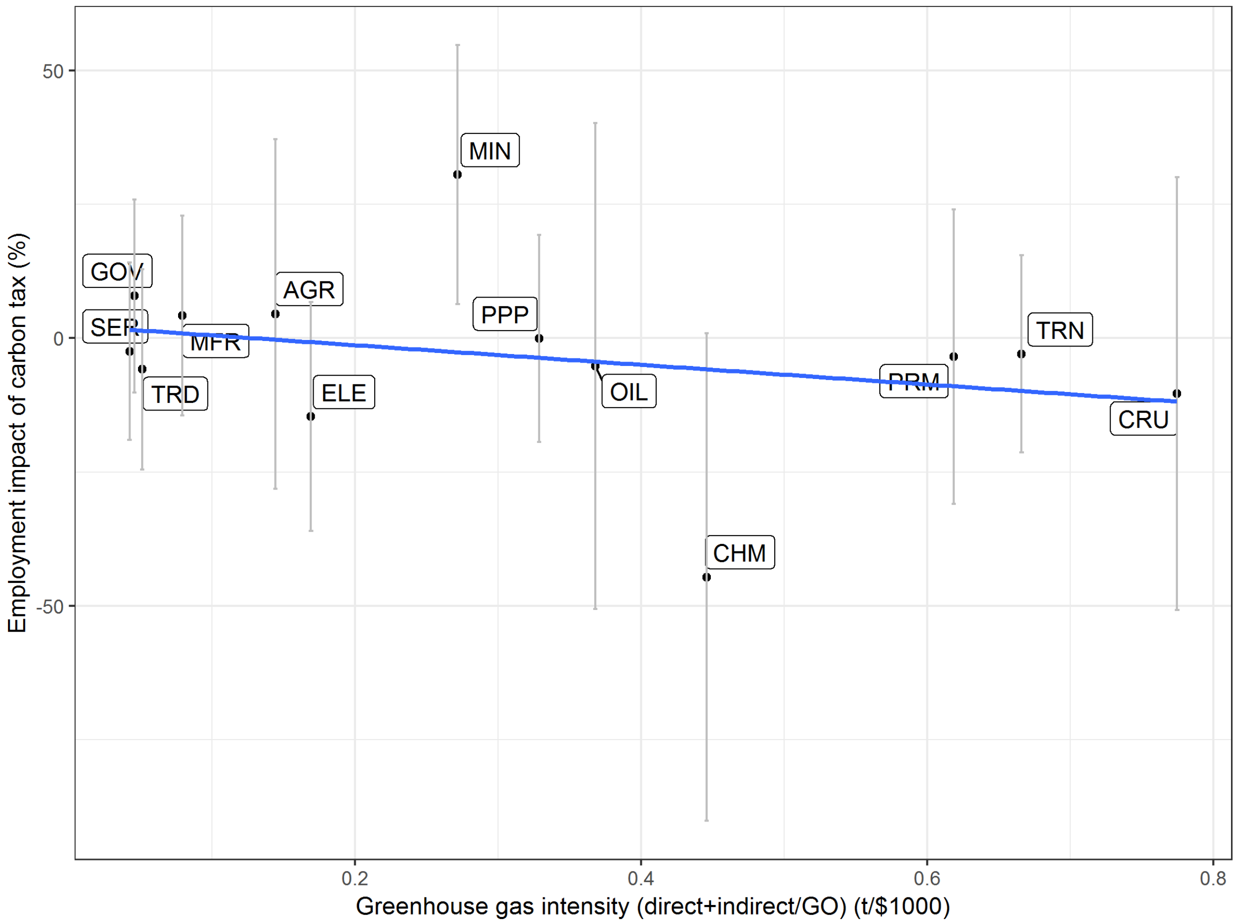

Figure 4: Results from alternative, non-parametric specification of the econometric model

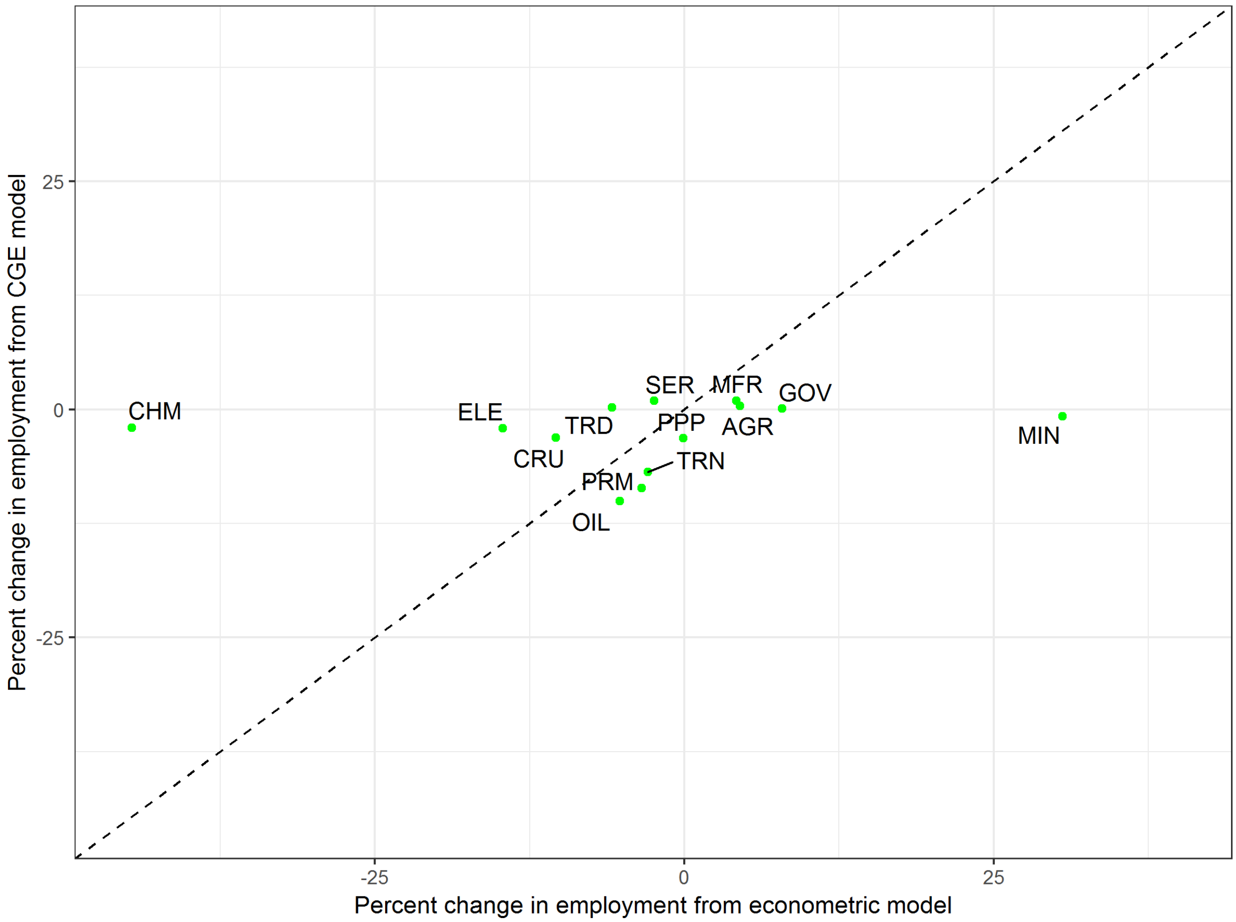

Figure 5: Comparison of econometric and CGE estimates of change in employment by sector associated with unilateral adoption of a \$30/t CO_2 tax in British Columbia. Non-parametric specification is used to generate the point estimates for the econometric model. Dashed black line has intercept of 0 and slope of 1; points would fall along this line if the two models were identical.

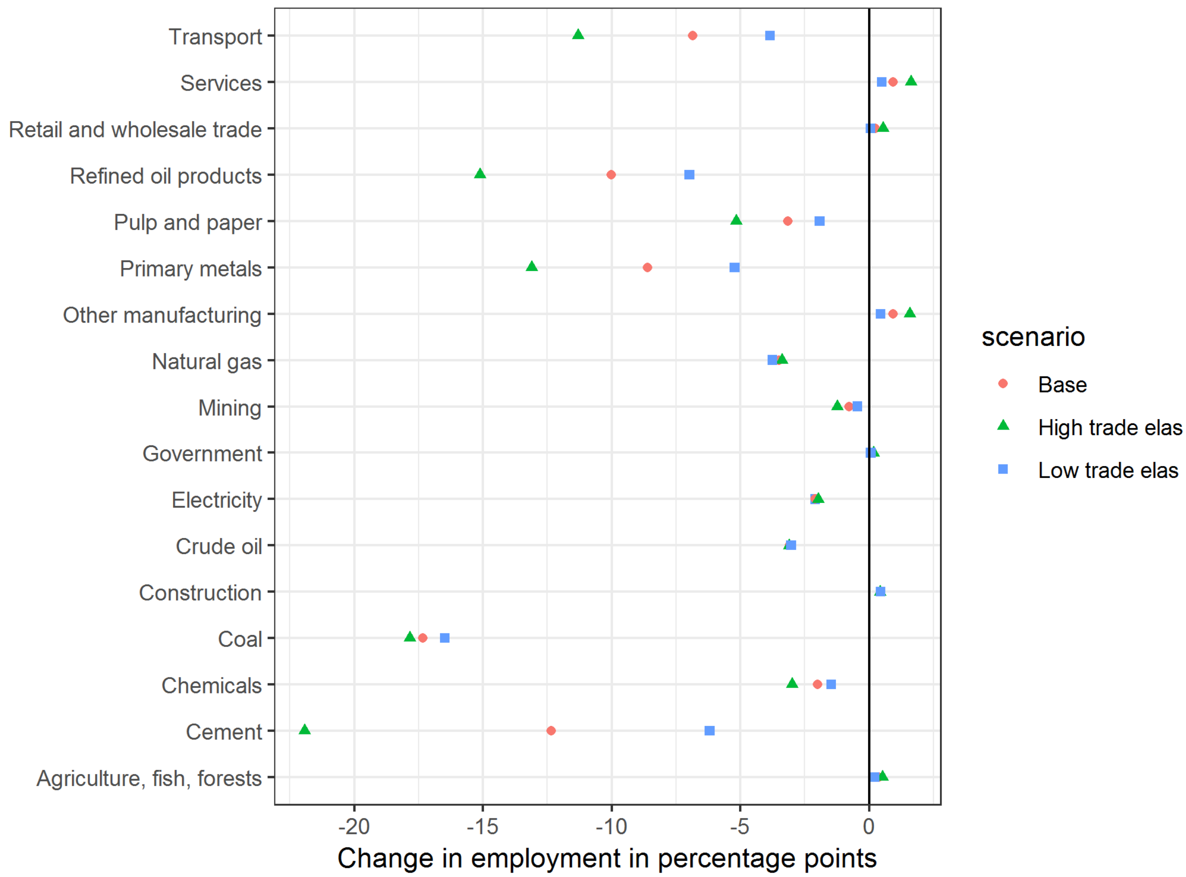

Figure 6: Sensitivity of CGE model results to differences in trade elasticity. In the high trade elas scenario, we double all trade elasticities; in the low trade elas scenario, we halve all trade elasticities.

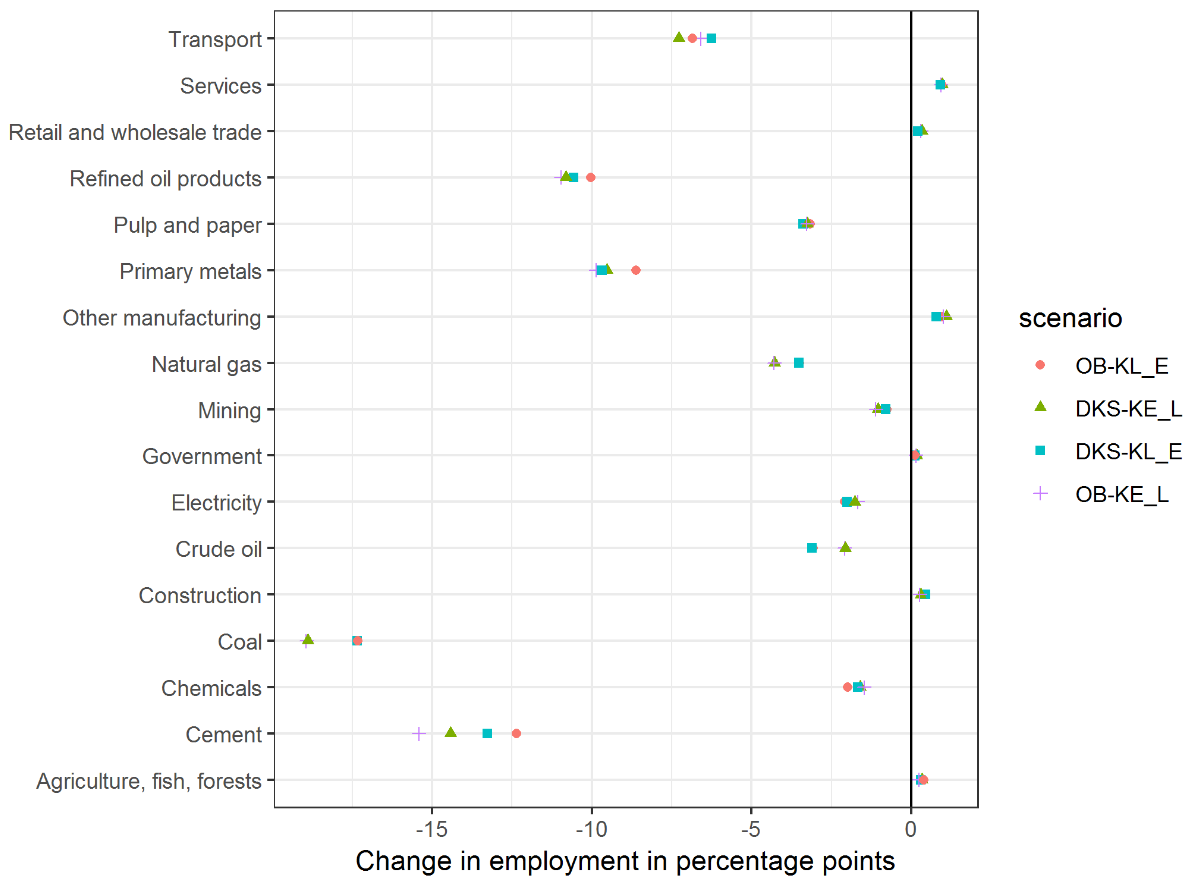

Figure 7: Sensitivity of CGE model results to differences in production nesting structure and elasticities. OB-KL_E is the base specification, taken from Okagawa and Ban 2008, where capital and labour are aggregated in a nest, and this nest is subsequently aggregated with energy. Points with a prefix of DKS are derived from Dissou, Karnizova, and Sun 2015.

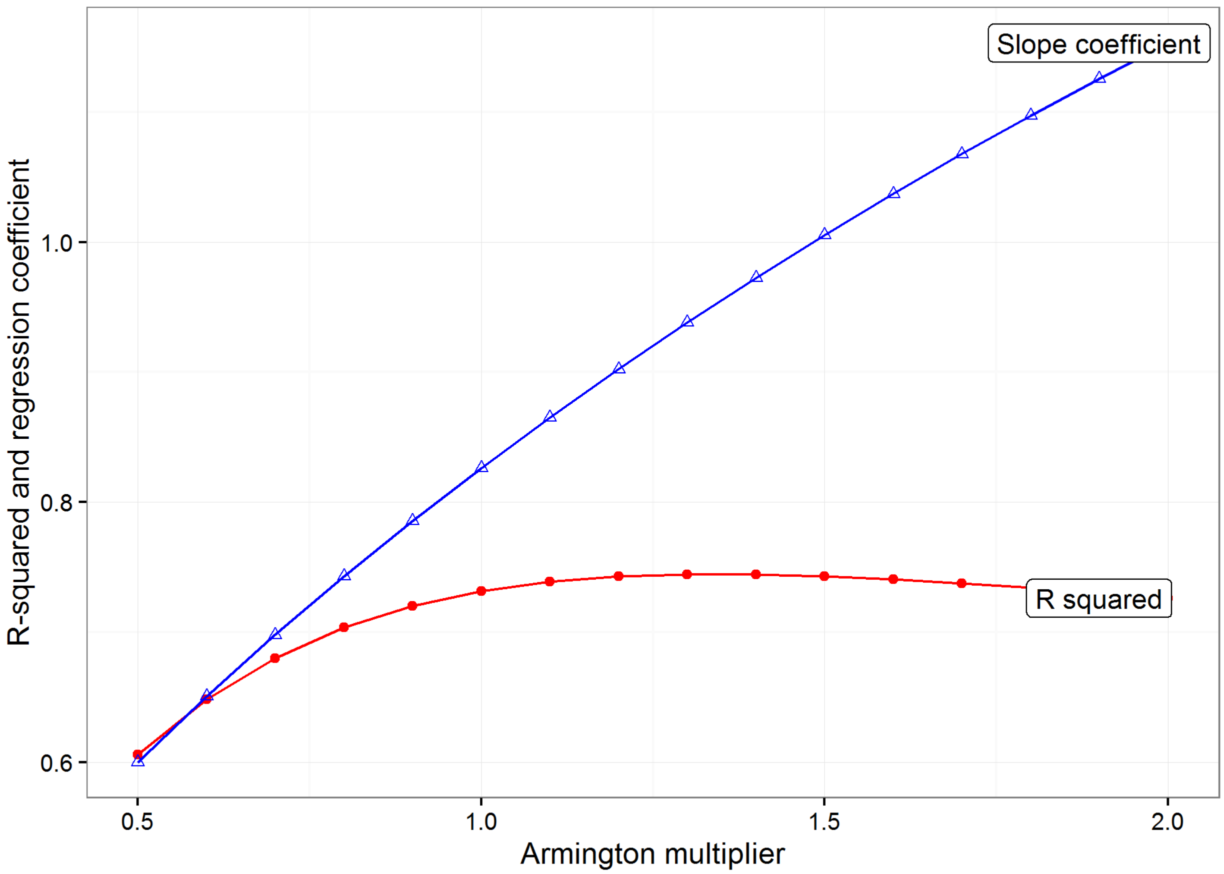

Figure 8: Comparison of CGE and econometric results with alternative values for trade elasticities.

7 Tables

Table 1: Benchmark sector profiles for British Columbia

Mnenomic	Description	Value of output ($B)	Value of exports ($B)	Emissions (Mt CO_2e)	Emission intensity (t/$1,000)	Export intensity
GAS	Natural gas	7.85	5.82	3.90	0.50	0.74
CRU	Crude oil	0.68	0.48	0.51	0.78	0.70
COL	Coal	1.16	1.07	1.20	1.05	0.92
OIL	Refined oil products	2.23	1.07	0.64	0.37	0.48
ELE	Electricity	4.35	0.91	0.66	0.17	0.21
AGR	Agriculture, fish, forests	8.20	2.37	0.86	0.14	0.29
MIN	Mining	2.76	1.77	0.63	0.27	0.64
CON	Construction	36.36	0.01	0.34	0.04	0.00
PPP	Pulp and paper	4.13	2.30	1.11	0.33	0.56
PRM	Primary metals	1.60	1.03	0.87	0.62	0.65
CHM	Chemicals	0.78	0.52	0.31	0.45	0.66
CEM	Cement	0.31	0.12	0.32	1.19	0.38
MFR	Other manufacturing	35.44	20.4	1.21	0.08	0.58
TRD	Retail and wholesale trade	35.36	2.59	0.63	0.05	0.07
TRN	Transport	30.75	10.2	12.13	0.67	0.33
SER	Services	138.18	19.0	1.46	0.04	0.14
GOV	Government	42.10	0.70	0.73	0.05	0.02

Notes: Emission intensity is calculated as the direct plus indirect emissions subject to the carbon tax (i.e., due to fossil fuel combustion and originating in British Columbia) divided by gross output. Export intensity is exports divided by gross output.

Table 2: Regressions of log employment on carbon-tax variables and fixed effects

lnL	(1)	(2)	(3)	(4)
$EI_{ijr} \times \tau_{rt}$	-0.0053**	-0.0053**	-0.0052**	-0.0052**
	(0.0025)	(0.0022)	(0.0022)	(0.0022)
τ_{rt}	0.0018	0.0052**	0.0011	
	(0.0013)	(0.0028)	(0.0024)	
N	4,181	4,181	4,181	4,181
R^2	0.57	0.61	0.61	0.61
Industry × province FE	Y	Y	Y	Y
Industry × time FE	Y	Y	Y	Y
Province × time FE				Y
Province trends			Y	
Province trends^2			Y	
Province				
× Oil price		Y	Y	
× US unemployment		Y	Y	
× World export		Y	Y	
Population by age		Y	Y	

Notes: Dependent variable is log of employment. EI_{ijr} is emission intensity level for industry i of sector j from reion r; τ_{rt} is a carbon tax variable (i.e., for BC industries 0 if t = pre-carbon tax period, 10 if $t = 2008$, 15 if $t = 2009$, ..., 30 if $t = 2012$ and 0 for ROC industries.). Standard errors are clustered by sector × province, reported in parentheses. Oil is crude oil price (WTI), US unemployment is US unemployment rate, and World export is world export volume relative to world GDP (constant at 1913). These three control variables are interacted with provincial dummies. Population by age breaks up annual total provincial population in sevel categories: kids (0 to 14 years), teens (15 to 19 years), adults in 20s, adults in 30s, adults in 40s, adults in 50s, and seniors (above 60 years). *** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.

Table 3: Regressions of CGE point estimates on econometric point estimates

	CGE simulation	
	Unweighted	Weighted
	(1)	(2)
Econometric prediction	0.85***	0.81***
	(0.13)	(0.07)
Constant	-1.48^{*}	-1.34^{***}
	(0.79)	(0.24)
N	17	17
R^2	0.73	0.89

Notes: Dependent variable is predicted change in sector employment with a $30/t carbon tax from the CGE model. Independent variable is predicted change in sector employment with the same policy from the preferred specification of the econometric model. *** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.

Table 4: Pseudo-regressions using data generated from CGE model.

lnL	(1)	(2)	(3)
$EI_{jr} \times \tau_{rt}$	-0.0046***	-0.0044***	-0.0045***
	(0.0006)	(0.0006)	(0.005)
$TI_{jr} \times \tau_{rt}$			0.0002
			(0.0005)
τ_{rt}	0.0004**	0.0004*	0.0004*
	(0.0002)	(0.0002)	(0.0002)
N	888	888	888
R^2	0.85	0.85	0.85

Notes: Dependent variable is log of employment for sector j in province r at time t. EI_{jr} is emission intensity level for sector j from reion r and TI_{jr} is the share of output that is exported in sector j in province r in 2007. All specifications includes sector $\times$ time fixed effects and sector $\times$ province fixed effects. Standard errors clustered by province $\times$ sector are in parentheses. *** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.

Table 5: Regression robustness check controlling for trade-exposure

	(1)	(2)
$EI_{ijr} \times \tau_{rt}$	-0.0052**	-0.0051**
	(0.0022)	(0.0022)
$TI_{ijr} \times \tau_{rt}$		0.0012
		(0.003)
τ_{rt}	0.0011	0.0003
	(0.0028)	(0.0033)
N	4,181	4,181
R^2	0.61	0.61

Notes: Dependent variable is log of employment for industry i of sector j in province r at time t. TI_{ijr} is defined as the share of output that is exported in industry i of sector j in province r in 2007. All specifications include the same specification as column (3) in Table 2. Standard errors clustered by sector $\times$ province are in parentheses. *** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.