

The Efficiency and Distributional Effects of Alternative Residential Electricity Rate Designs

Scott P. Burger,^a Christopher R. Knittel,^b Ignacio J. Pérez-Arriaga,^c Ian Schneider,^d
and Frederik vom Scheidt^e

Residential electricity tariffs typically distort—and thus do not allow consumers to respond to—the marginal cost of energy consumption. Rates are typically constant across time and location, despite the fact that short-run marginal costs can vary dramatically. As of the end of 2016, less than one quarter of one percent of residential customers in the U.S. faced electricity prices that reflected the real-time marginal cost of energy production. Furthermore, the bulk of system costs are recovered through volumetric charges—that is, charges per-unit of energy consumed—despite the fact that a substantial fraction of these costs are fixed in the short term. More economically efficient rate designs could substantially improve market efficiency.

This reality has led many regulators, policy makers, consumer advocates, and utilities to call for improved tariff designs. For example, the New York Department of Public Service recently called for “more precise price signals... that will, over time, convey increasingly granular system value.” New York is not an anomaly. In 2017, regulators in 45 of 50 U.S. states and the District of Columbia opened dockets related to tariff design or made changes to tariff design. Similarly, in November 2016, the European Commission issued a sweeping set of rulings, with tariff design as a centerpiece.

The economic pressure to redesign electricity rates is countered in part by concerns among policy makers and regulators of how more efficient rate structures might impact different socio-economic groups in terms of both average bills and bill volatility. For example, the Massachusetts Department of Public Utilities (MADPU), the New York Department of Public Service (NYDPS), and the California Public Utilities Commission (CPUC) all list concerns about the distributional impacts of rates in their principles for rate design. Regulatory decisions highlight these concerns: in the U.S. in the second quarter of 2018, state electricity regulators rejected over 80% of utility requests to increase fixed charges, frequently citing the potential impacts on low-income customers.

This paper examines the distributional and economic efficiency implications of residential electricity tariffs using interval metering data—measuring electricity consumption every 30 minutes—for more than 100,000 customers in the Chicago, Illinois area. Our work leads us to a number of novel findings. First, we find that, holding the proportion of fixed and volumetric charges in the tariff constant, movements towards more time-varying rates tend to decrease annual electricity expenditures for low-income customers. However, increases in customer fixed charges tend to increase expenditures for low-income customers who, on average, consume less electricity than their more affluent counterparts. The net effect of instituting a rate design with real-time energy prices and uniform fixed charges for residual cost recovery is a monotonic negative relationship between in-

a PhD Candidate, Institute for Data, Systems, and Society, MIT. Email: sburger@mit.edu.

b Shultz Professor of Applied Economics, Sloan School of Management, Director of the Center for Energy and Environmental Policy Research, Co-Director, Electric Power Systems Low Carbon Energy Center, MIT and NBER. Email: knittel@mit.edu.

c Visiting Professor, Sloan School of Management, MIT. Professor and Director of the BP Chair on Sustainable Development, Comillas University. Director of Energy Training, Florence School of Regulation. Email: ipa@mit.edu.

d PhD Candidate, Institute for Data, Systems, and Society, MIT. Email: ischneid@mit.edu.

e Institute of Information Systems and Marketing, Karlsruhe Institute of Technology. Email: frederik.scheidt@kit.edu.

come and changes in expenditures (that is, as income increases, changes in expenditures decrease). Second, in our sample, the economic distortions of recovering residual network and policy costs through volumetric tariffs likely outweigh the distortions that emerge from charging an energy price that does not reflect the underlying time- and location-varying cost of energy. Finally, we find that changes to fixed charge designs can preserve the efficiency gains of transitioning to efficient residual cost recovery while mitigating undesirable distributional impacts. We highlight three methods for designing fixed-charges for residual cost recovery—based on customer demand characteristics, income, or geography—that mitigate the regressiveness of fixed charges.

The Impact of a Carbon Tax on the CO₂ Emissions Reduction of Wind

Chi Kong Chyong,^a Bowei Guo,^b and David Newbery^c

Energy policy aims to reduce emissions at least long-run cost while ensuring reliability. Policies to support wind or solar PV, improve efficiency, or shift peak demand need to be assessed on the cost of the emissions reduced. Ofgem (2018) in its *State of the market 2018* is a good example, comparing the cost effectiveness of various UK energy policies. This paper shows how to estimate CO₂ reductions in electricity from wind deployment in both the short and long run. Both need to be estimated to adequately measure carbon savings.

Clearly, just how much wind and other renewables reduce emissions will depend on the carbon price and the types of plant on the system. The carbon price will affect the merit order and influence which plant type runs, and hence on the fuel and carbon that wind displaces. The main EU instrument for setting the price of CO₂ is the Emissions Trading System (ETS). From plausible levels of the EU Allowance (EUA) price of €20-30/tonne in 2008, the price sharply declined after the *Renewables Directive* reduced demand without withdrawing EUAs, and for long periods of time after 2012 has remained well below €10/tonne, at which level it had little effect on the fuel mix.

The UK Government has committed to tough carbon limits in the *UK Climate Change Act 2008* and introduced a Carbon Price Floor (CPF) in the 2011 Budget that applies to fossil fuels used to generate electricity. To implement the CPF, the Treasury publishes the Carbon Price Support (CPS, a carbon tax in addition to the EUA price) based on forward EUA prices at the time of the autumn budgets to come into effect at the start of the fiscal year in the following April. The CPS was revised several times to bring the total CO₂ price up to the announced CPF trajectory that was planned to reach £₂₀₁₁30/tonne by 2020 and £₂₀₁₁70/t by 2030. After the failure of other EU countries to either reform the ETS or impose a similar CPF, the Government froze the CPS in 2016 at £18/t until 2021.

The impact of this considerable increase in the cost of fossil fuel for electricity generation has been dramatic. Before the introduction of the CPS coal generation was cheaper than the most efficient Combined Cycle Gas Turbines (CCGTs), so that gas was the marginal fuel in the mid-merit part of the market. After the introduction of the CPS coal eventually became the most expensive fossil fuel (in April 2015, when the CPS increased from £9.55/tCO₂ to £18.08/tCO₂), starting a massive switch from coal-fired generation to gas. The share of coal fell from 41% in 2013 to 6% in 2018.

a EPRG, University of Cambridge Judge Business School, Cambridge, UK. E-mail: k.chyong@jbs.cam.ac.uk.

b Faculty of Economics, University of Cambridge. E-mail: bg347@cam.ac.uk.

c Corresponding author. EPRG, Faculty of Economics, University of Cambridge. E-mail: dmgn@cam.ac.uk

Great Britain (the CPS does not apply in Northern Ireland) therefore offers an excellent test-bed for the impact of a carbon tax (the CPS).

Wind is hard to forecast with much accuracy day-ahead when the time comes to decide which types of generation to commit and run. As wind varies from moment to moment, the carbon displaced will depend on the plant operating and its flexibility. We study this short-run impact econometrically to find the main drivers of the short-run displacement achieved.

Policies are chosen for their long-run impact. Governments set targets for the future share of renewable electricity and carbon budgets. These policies will affect the future fuel mix, and hence the dispatchable plant available daily. We determine this long-run impact with a unit commitment dispatch model of the 2015 GB electricity system. We examine the effect of increasing wind capacity by varying amounts up to 25%. Long run has the conventional meaning that it is a period over which wind capacity can change, in contrast to the short run in which wind capacity is fixed but its output varies. However, in our model we hold the conventional plant mix and interconnector flows constant as we vary carbon prices and wind capacity to isolate wind capacity's impact.

The econometric models estimate the marginal (coal and gas) displacement of wind (in MWh coal or gas/MWh of wind). Given this, we estimate the short-run marginal displacement factor (SR-MDF) of wind on emissions, tonnes CO₂/MWh wind. There are two main advantages for this approach: first, it explains the underlying mechanisms that drive the dynamics of the MDF; second, it allows us to study the counterfactual, i.e. what if the CPS is not implemented or what if it were higher? We are unable to deliver the counterfactual without knowing the underlying drivers. The econometrics also estimates the impact of demand changes on emissions as the Marginal Emissions Factor (MEF) in tonnes CO₂/MWh.

Results

A simple examination of the evolution of the fuel mix over time strongly suggests that gas has displaced coal, and that wind has displaced both, but as the clean spark and dark spreads have varied substantially over time with varying fuel and carbon prices, a more detailed examination was undertaken to tease out the various effects.

The short-run impact of varying wind half-hourly on the fuel mix and emissions was explored econometrically, using changes in fuel and carbon prices as well as in wind output and final demand over 2012-17 to create sufficient variation to identify the drivers of the half-hourly marginal displacement of wind. The econometric study suggests that the short-run MDF of wind depends on the level of demand (i.e. which fuel type is running at the margin), the merit order, and the flexibility of fossil plants. When demand is low, base load plant responds more strongly to short-run wind changes. However, when demand is high and so is its variability, more flexible CCGTs are preferred for responding to wind changes. Because of this, CCGTs dominate as the (short-run) marginal fuel during peak hours (07:00-23:00) regardless of the merit order, while coal would only be the (short-run) marginal fuel during off-peak hours (23:00-07:00) when coal is cheap enough to run as base load. Consequently, as the CPS switches the merit order, coal becomes the mid-merit fuel but not the flexible marginal fuel, making CCGTs the marginal fuel for the entire day. Because of the much lower carbon intensity of gas, the SR-MDF decreases from 0.44 tCO₂/MWh in the counterfactual absence of the CPS in our reference year, 2015, to 0.41 tCO₂/MWh with the CPS. The EU Market Stability Reserve sharply raised carbon prices in 2018, and in our counterfactual of a total GB carbon price of £37/tCO₂, the SR-MDF falls further to 0.36 tCO₂/MWh. In each case the MEF is either the same or slightly higher than the SR-MDF.

The unit commitment simulation model is able to explore the effect of different total carbon prices on the carbon savings from a significant but modest increase (25% of its 2015 level) in installed wind capacity, holding fuel prices, fossil capacity and interconnector trades constant. This showed that with 2015 gas and coal prices, introducing the CPS as an additional £18/tCO₂ on an EUA price of £6/tCO₂ switched coal from base-load to mid merit, making it somewhat more responsive to changes in wind capacity, slightly increasing the long-run carbon benefits of more wind. The LR-MDF rises from 0.51 tCO₂/MWh with no CPS to 0.60 tCO₂/MWh with the CPS. Increasing the total carbon price further, however, leads to a decrease in coal generation as it moves more to peak hours, resulting in a smaller decrease in emissions with increases in wind capacity (because coal plants are either running at minimum stable generation or not at all) to a LR-MDF of 0.57 tCO₂/MWh. The LR-MDF is systematically higher (by 15% with no CPS to 46% with the CPS) as the unit commitment model assumes perfectly predictable future wind and demand and can therefore schedule more of the less flexible and higher emitting coal plant.

Both the simulation and the econometrics confirm that the impact of wind depends quite sensitively on the state of the system—which plant are running and whether they are constrained by minimum loads, capacity, or ramping limits. That in turn depends on fuel and carbon prices and the levels of residual demand. Different countries have very different plant mixes, and so the carbon benefits of additional renewables capacity will also vary, while over time, fuel and carbon prices as well as the plant mix will also vary. This paper shows how the emissions benefits can be measured for a given plant mix and set of fuel and carbon prices, implying that country level detailed modeling will be needed to understand their impacts. The long-run impact of increasing wind capacity is largely captured by the LR-MDF, although there will be some short-run variability that will, at least with our estimates based on GB conditions, somewhat reduce these values as more flexible gas is required to address unpredicted fluctuations in residual demand.

Effects of Ownership and Business Portfolio on Production in the Oil and Gas Industry

Binlei Gong^a

The Shale Revolution and high volatility in oil prices (especially in 2008 and 2014) have overwhelmingly reconstructed the petroleum industry in the last decade. These events not only significantly affected productivity but also overwhelmingly reshaped the production process of the industry. As a result, the shape of the production function may be time-variant. Moreover, heterogeneity across companies is also a big concern as many multi-product (oil and gas) and multi-segment (upstream and downstream) firms exist, both state-owned and privately-owned. Since different segments and products involve different technologies, and companies with different ownerships are found to be completely distinct, each has its own production process and input–output relation.

To summarize, both ownership and business portfolio, along with time, can influence the production function. The first puzzle is that the classic method assumes the shape of the production function is fixed, which fails to capture the varying input–output relation over time and across firms. To solve this issue, this paper employs a varying production function to better estimate the true data-generating process, where the production function not only can shift, as in the classic model,

^a Corresponding author: Department of Agricultural Economics and Management, School of Public Affairs, China Academy for Rural Development (CARD), Zhejiang University, Hangzhou 310058, China. E-mail: gongbinlei@zju.edu.cn

but can also be reshaped depending on the aforementioned factors. Although the varying coefficient model captures the heterogeneity in time, ownership, and business portfolio, their effects are modeled nonparametrically and are therefore hard to clarify. This second problem motivates us to further decompose input elasticities and total factor productivity using a second-step parametric regression, so that their impacts can be better interpreted.

Three central contributions are made in this paper. First, a varying coefficient method is used to better capture the fundamental transition of the petroleum industry over the past decade. Second, this study extends the comparison between National Oil Companies (NOCs) and privately-owned International Oil Companies (IOCs) in literature by not only interpreting more accurate effect of ownership on productivity than was estimated with bias, but also investigating its impacts on the input–output relation that was overlooked. Third, this paper also controls the heterogeneity in segment-wide and product-wide business portfolios and predicts their effects on production, which contributes to the studies of multi-product, multi-segment and multidivisional firms.

Using panel data from 110 of the largest oil and gas companies from 2009–2016, the empirical results show that labor elasticity was increasing, elasticities of oil and gas reserves were decreasing, and refinery elasticity was relatively flat. The changes in the productivity of these petroleum giants follows the business cycle, which achieved tremendous growth after the 2007–2009 financial crisis but lost momentum following the oil price crash since the end of 2014. Furthermore, segment-wide and product-wide portfolios, as well as ownership, significantly affect the input elasticities and therefore the shape of the production frontier, which in turn confirms the importance and necessity of employing the varying coefficient model, rather than the classic model. Finally, IOCs, gas production and downstream activities, *ceteris paribus*, are more productive than NOCs, oil production and upstream activities, respectively.

The empirical results deliver the following policy implications: 1) the government should be very careful in putting pressure on NOCs to fulfill their non-commercial objectives; 2) privatizing NOCs may be a good option to prevent shortage in oil and gas supply and higher energy prices; and 3) oil companies can hire more oil field service companies to explore and produce oil and gas.

The Effects of Fuel Prices, Environmental Regulations, and Other Factors on U.S. Coal Production, 2008-2016

John Coglianese,^a Todd D. Gerarden,^b and James H. Stock^c

From 2008 to 2016, U.S. coal production fell 37 percent, from 1,172 million short tons to 739 million short tons. This decline in coal production drove a 23 percent reduction in power sector CO₂ emissions, as well as providing health co-benefits through reduced particulates and other local pollutants, but was also associated with substantial job losses, with employment in coal mining falling from 87,000 in 2008 to 52,000 in 2016.

Public discussions of the reasons for this decline focus on three explanations: the decline in natural gas prices as a result of fracking, environmental regulations affecting coal-fired power

a Research & Statistics Division, Board of Governors of the Federal Reserve System.

b Corresponding author: Charles H. Dyson School of Applied Economics and Management, Cornell University, Ithaca, NY 14853, USA. Email: gerarden@cornell.edu.

c Department of Economics and Harvard Kennedy School, Harvard University, and National Bureau of Economic Research.

plants, and the role of state-level renewable mandates on increasing wind and solar generation. Other factors are (small) changes in total electricity demand, changes in demand for metallurgical coal, and changes in export demand.

We decompose the decline of coal production from 2008 to 2016 into nine components: changes in the relative price of coal to gas, environmental regulations, renewable portfolio standards, total electricity demand, improvements in power plant heat rates, demand for metallurgical coal, net exports, industrial demand, and an unexplained residual. To do so, we use a combination of econometric and accounting methods. The econometric analysis uses state-level data on monthly electricity generation by coal, gas, and other sources, prices of gas and coal, and environmental and renewable mandates to estimate their various effects on the generation share of coal-fired power in total electricity generation. The remaining six factors are incorporated using an accounting identity approach. One environmental rule, the Mercury and Air Toxics Standard (MATS), took effect nationwide so is susceptible to neither an econometric nor accounting approach. One method for complying with the MATS rule is closing a plant. We therefore perform a differences-in-differences analysis of the effect of the rule on planned plant closures and add the results to the econometric decomposition.

According to our estimates, the declining price of natural gas relative to coal, on an energy-adjusted basis, explains 92 percent (SE = 2.5 percentage points) of the decline in coal production, or approximately 397 of the 433-million-ton decline. An additional six percent (SE = 2.2 percentage points) of the decline is explained by environmental regulations, primarily the Cross-State Air Pollution Rule and the MATS rule. We attribute a small amount of the decline, two percent (nine million tons), to the adoption of renewable portfolio standards. The remaining six factors contribute small, largely offsetting amounts to the change in total coal production.

These results underscore that the dramatic shift away from coal during this period was primarily a result of market forces, not government policy. Looking ahead, they suggest that, absent additional policy, the prospects for a coal rebound, or for further declines, are intimately linked to the price of gas relative to coal. If that relative price remains stable at its currently low levels, one would not expect a rebound in coal, but neither would one expect a continued decline. Instead, continued reduction in CO₂ emissions from coal must stem either from future policy choices, from declining prices of solar and wind generation, or both.

Consumer Inattention, Heuristic Thinking and the Role of Energy Labels

Mark A. Andor,^a Andreas Gerster,^b Stephan Sommer^c

Research has shown that buyers of energy-using durables tend to pay less attention to opaque lifetime cost than to salient purchasing prices. Indeed, inattention to energy cost is one aspect that may explain the low tendency of consumers to invest in cost-effective efficiency technologies, i.e. the so-called “energy efficiency gap”. To bridge this gap, policy-makers all around

a Corresponding author. RWI—Leibniz-Institute for Economic Research, Hohenzollernstr. 1-3, 45128 Essen. E-mail: andor@rwi-essen.de

b University of Mannheim, L 7, 3-5, 68161. E-mail: gerster@uni-mannheim.de.

c RWI—Leibniz-Institute for Economic Research, Hohenzollernstr. 1-3, 45128 Essen. E-mail: sommer@rwi-essen.de.

the world have introduced energy labels on electric appliances. Some of them summarize energy information in an intuitive way, such as energy efficiency classes, which consumers might employ instead of the more detailed information on energy use.

This paper analyzes the potential of energy labels with energy efficiency classes to influence consumer choices. We propose a conceptual model of investment decisions on energy efficiency that explicitly takes into account consumers' inattention to operating cost and their tendency to employ decision heuristics based on energy efficiency classes. In our empirical analysis, we first investigate whether households have a willingness-to-pay for energy efficiency class differences per se, i.e. irrespective of energy use differences. Second, we analyze how an increase in the salience of annual operating cost as well as the number of stimuli that compete for attention affect appliance choices. Third, we investigate the channels through which changes in the salience of operating cost and the number of competing stimuli operate.

We conduct a stated-choice experiment among 5,000 households and frame it as a purchasing decision on refrigerators that display the EU energy label and assign participants randomly into three groups. Participants in a control condition receive information on the appliances based on a label that only displays the annual energy use and energy efficiency classes. In the first treatment condition, we increase the salience of the cost component by adding estimated annual operating cost on the label. In the second treatment condition, participants see further non-energy related appliance characteristics that act as additional stimuli competing for attention.

Our results demonstrate that additional cost information on the EU energy label steers consumers to more energy-efficient appliances. In contrast, exposing consumers to additional non-energy-related appliance characteristics can impede the choice of energy-efficient appliances.

We further show that consumers value energy efficiency class differences per se: even when differences in energy use are marginal, two thirds of consumers are willing to pay at least 30 EUR for a better energy efficiency class. This holds particularly for individuals with a higher cognitive burden of decision making, for instance, respondents who are uninformed about electricity prices. The presence of such decision heuristics is far from innocuous, as they can distort consumers' valuation of energy efficiency with far-reaching implications on producers' innovation activities. In particular, it may discourage producers from developing new energy efficiency technologies that do not result in better energy efficiency classes.

Moreover, we find that adding information on operating cost to the EU energy label works through two distinct channels: it increases attention to operating cost and decreases the valuation of energy efficiency class differences. Hence, we identify a substitution effect between operating cost information and the information incorporated in energy efficiency classes. Consequently, raising the salience of operating cost can have ambiguous effects on the adoption of energy-efficient appliances: it increases households' attention to electricity use, but also decreases consumers' valuation of energy efficiency class differences. This finding enhances the understanding of how coarse summary information on labels interacts with more detailed information. This has implications for the design of not only energy labels, but also labels in other domains.

Based on our results, we argue that positive welfare effects of a label revision seem likely. First, the provision of operating cost can be considered a "pure nudge", i.e. a behavioral intervention that only informs previously uninformed consumers, but has no further effects. We do not expect private welfare of consumers to decrease after being better informed. Second, implementation cost to introduce and update the operating cost information seem limited. Finally, a revision of the label promises significant reductions in negative externalities associated with electricity consumption. As more than 15 million refrigerators and millions of other appliances are sold annu-

ally in the EU, even small improvements in the energy efficiency of the appliance stock can help to reduce electricity consumption and the associated carbon emissions substantially.

On the Viability of Energy Communities

Ibrahim Abada,^a Andreas Ehrenmann,^b and Xavier Lambin^c

The notion of energy communities has received increased interest over the past few years, fostered by better information and communication technologies and an increase in environmental awareness. Energy communities are seen as crucial for facilitating the decentralization of power production and enhancing the management of energy resources at a local level. Small energy communities are thus blossoming in many countries, often bolstered by substantial support from policy-makers (for example, in the form of feed-in tariffs). In March 2019, the European Commission concluded negotiations on the Clean Energy for All Europeans package, adopting new market design rules where consumers are put at the center of energy systems and citizens energy communities are promoted. However, despite the potential profits made by such communities, there is no guarantee that they will be viable. In the event a community would be created, a subset of participants may indeed find it profitable to exit the community and create another one of their own if not properly remunerated, which would therefore falsify the creation of the community in the first place.

This paper analyzes the stability of an energy community using the means of cooperative game theory. The basic idea is that if a community fails to provide stability, then by definition the community will not form in the first place. For that purpose, we consider a narrow definition of an energy community: households of a common building or close geographical area may decide to combine their effort and jointly build solar panels on their roofs (or windmills in a nearby field). Instead of individual meters, they can then decide to install only one and use it for the whole community. There is therefore one source of costs (the costs of installation of the renewable resource), and two sources of gains: aggregation gains, in the form of decreased network fees, and energy gains, as the renewable energy can be consumed at zero marginal cost or re-injected in the network and given a feed-in tariff. Depending on the structure of the community and the sharing rule that is used to allocate the gains among its members, we are able to determine whether the community can materialize in the long-term or not.

A first takeaway of our research is that the most basic sharing rules (per-capita, pro-rata of consumption or peak demand) usually fail to provide adequate remuneration to all players. In that case, some households may decide to opt out from the community. This finding casts strong doubts on the coherence between the concept of self-organized energy communities and the spirit of tariff equalization: in many countries, equal access to electricity both in terms of quality and price is a prerequisite. However, applying such equality concept to energy communities may render them unstable (meaning equivalently that they won't materialize in the first place). To promote communities, regulators may therefore want to question the extent to which tariff equality should be applied. More elaborate sharing rules, such as the Shapley value or the minimum variance allocation, though slightly more complex, have desirable properties and are more likely to enable communities to share their gains thereby enabling them to be viable. We acknowledge that there is a tradeoff between the

a Corresponding author. ENGIE, CEEME-Tractebel. E-mail: ibrahim.abada@engie.com.

b ENGIE, CEEME-Tractebel and Energy Policy Research Group, University of Cambridge.

c Grenoble Ecole de Management.

efficiency of the sharing rule we propose, and its complexity. Although we discuss some solutions to simplify the rules and make the resulting tariffs more readable, some consumers may not be willing to engage in relatively complex sharing schemes. When the community cannot be stable, the intervention of a local social planner or a change in network tariffs may be required to restore efficiency. If such an intervention is not desired, we propose a way to optimally split the whole energy community into smaller stable groups of consumers, so that the lost value when splitting is minimized.

The Impact of U.S. Supply Shocks on the Global Oil Price

Thomas S. Gundersen^a

The summer of 2014 marked the end of a prolonged period of remarkably high oil prices. From having the Brent crude oil price hover between \$110 and \$120 per barrel for almost 4 years, the oil price fell by more than 75 per cent to hit the \$28 mark in January 2016. There are several possible explanations for why prices decreased so dramatically during this period. Among the most popular are a weakening global economy that reduced the demand for crude oil, a strategy by OPEC to squeeze out high-cost producers by increasing output (and thus lowering prices) and an unprecedented surge in U.S. crude oil production caused by the ‘U.S. shale oil boom’.

It is the latter explanation that I explore in this paper. Particularly, to what extent, if at all, did the U.S. shale oil boom contribute to the oil price movements leading up to and beyond the decline of mid-2014?

In order to capture the effects of increased American crude oil production, I exploit that the U.S. until 2016 had a crude oil export ban in place. With the ban, the change in crude oil imports is an indirect measure of U.S. crude oil production with the advantage that it relates directly to oil prices abroad. A reduction in U.S. crude oil imports implies that there is a higher domestic availability that has a negative effect on oil prices outside the U.S.

By using a structural econometric model, I estimate the effects of supply and demand shocks on the oil price. Importantly, I get an estimate of the effect of negative shocks to U.S. crude oil imports on the price of imported crude oil.

The results show strong support for the U.S. shale oil boom affecting oil prices negatively. Following a U.S. import shock that reduces U.S. imports by 1 per cent, the oil price falls with almost 2 per cent. The relative contribution of these shocks on oil price variation is 13 per cent over the 2003-2015 period. However, the analysis also shows that during the years leading up to 2014, the cumulative effect of these shocks did not push the oil price downwards until late 2013 despite crude oil production in the U.S. growing rapidly from 2011. Supplementary information from industry sources suggests that this delay was due to insufficient pipeline capacity and incompatible refining equipment in the downstream supply chain hindering shale oil from being adopted by the U.S. refining sector and thus maintaining the reliance on foreign imports. These frictions are most clearly observed as the heavy discount put on the WTI oil price relative to the Brent crude. Finally, I calculate that the oil price would have been \$10 higher in the absence of U.S. import shocks from 2014 onwards.

The results put forward in this paper suggest that the United States to a larger extent than before is able to influence oil prices globally—perhaps for the first time since the first OPEC shocks

^a Centre for Applied Macro- and Commodity Prices (CAMP), Department of Economics, Nydalsveien 37, N-0484 Oslo, Norway. E-mail: thomas.s.gundersen@bi.no

in the 1970s. Further, the findings also hint to that future oil price declines stemming from the U.S. might be predicted by a narrowing of the spread between the Brent crude and the WTI prices, implying that bottlenecks in the domestic supply chain is reduced and imports of foreign oil can be displaced to a larger extent.

Did U.S. Consumers Respond to the 2014–2015 Oil Price Shock? Evidence from the Consumer Expenditure Survey

Patrick Alexander^a and Louis Poirier^b

The oil price decline of 2014–2015 was both unanticipated and large in magnitude, leading to expectations of a significant impact for the U.S. economy. Since then, there continues to be debate over how much this shock affected U.S. GDP, and whether increased consumer spending due to gasoline price savings was an important channel here.

In this paper, we examine whether U.S. households who were particularly exposed to the oil price shock changed their consumption behavior after the shock, using representative micro data from the Consumer Expenditure Survey from 2013–2015. We propose a difference-in-difference identification strategy by comparing consumption responses of a treatment group of households that owned a vehicle with responses of a control group that did not own a vehicle. For robustness, we also compare a treatment group of households that are above the 20th percentile in gasoline reliance with a control group of households below the 20th percentile in gasoline reliance. Our study is the first to examine this topic using representative micro data for U.S. households that covers spending on all types of products.

Our findings suggest that households that were exposed to oil prices significantly increased their spending after the oil price shock. In terms of magnitude, we find that the marginal propensity to consume (MPC) out of gasoline price savings was greater than 1 during this period, which is larger than estimates found in several other studies that examine this episode.

Across products, we find that exposed households increased their consumption of non-essential items, including alcohol, apparel, entertainment, and food away from home. In addition, we find that these households especially increased their expenditure on transportation products (e.g. motor vehicles), which are complementary to oil. That transportation products tend to be items where expenditures are large and discrete helps to explain why we find that the MPC out of gasoline price savings was above 1.

Overall, our findings suggest that the oil price decline of 2014–2015 had very significant positive effects on U.S. consumer spending.

a Corresponding author: Senior Economist, International Economic Analysis Department, Bank of Canada, 234 Wellington Street West, Ottawa, Ontario, Canada, K1A 0G9. E-mail: palexander@bankofcanada.ca.

b Corresponding author: Economist, International Trade and Finance Branch, Department of Finance, Government of Canada, 140 O'Connor Street #19, Ottawa, ON K1A 0G5. E-mail: louis.poirier@canada.ca.

Pipeline Capacity Rationing and Crude Oil Price Differentials: The Case of Western Canada

W.D. Walls^a and Xiaoli Zheng^b

This paper quantifies the impact of pipeline constraints on the crude oil price differentials between U.S. and Canadian markets. The quantitative analysis is crucial to understand the cost of inadequate pipeline capacity and the benefits of building new pipelines or expanding existing pipeline infrastructure.

Canada is a net exporter of crude oil and its major export destination is the U.S. In 2017, more than 95% of Canadian crude oil was transported from western Canada to U.S. markets through pipelines. Because Canada has limited refinery capacity, crude oil from western Canada is refined in the U.S. regional markets, particularly the Gulf Coast, where more than 60% of North American refinery capacity is clustered. Canadian crude oil is sold at a discount relative to the U.S. benchmark oil, West Texas Intermediate (WTI). One reason for this discount is that Canadian crude oil is lower quality than WTI. However, quality alone cannot explain the substantial price difference between WTI and Canadian crude oil. There is almost no quality difference between Mexican Mayan and the Canadian benchmark crude oil, Western Canadian Select (WCS). Nevertheless, Mayan is traded at much higher prices than WCS in U.S. refining hubs. Both Canadian and Mexican crude oils are shipped to refiners in the U.S. Gulf Coast. While Mexican oil is shipped by sea, the transportation of Canadian crude oil relies mainly on pipelines. Building new pipelines or expanding existing pipeline capacity in Canada is constrained. If pipeline capacity is insufficient, producers will face an oil glut, resulting in a low crude oil price relative to the prices of the rest of the world. However, the extent to which pipeline capacity constraints have increased price differentials between U.S. and Canadian oil markets remains an open question.

We exploit a unique dataset in Canada which directly represents the extent to which crude oil shipping nominations are excluded from pipeline service due to pipeline capacity constraints. Compared to previous studies, our measure of the pipeline capacity constraint has two properties. First, it directly reflects the quantity of crude oil that has been excluded from pipeline service due to insufficient capacity. Second, the metric of the capacity constraint is expressed in terms of the reduced percentage of total nominations. Using this measure, the degree of pipeline capacity constraint is comparable among different pipelines. Using the reduced percentage normalizes the degree of pipeline capacity constraint by the quantity of pipelines' shipping nominations. Our empirical analysis relates the measure of pipeline transportation capacity constraints to the discount of western Canadian crude oil relative to U.S. crude oils.

We estimate that the price differential between the U.S. markets and western Canadian increases by 3.6% for 1% increase in the degree of pipeline capacity constraints. Pipeline constraints in Canada have led to an average loss of as much as \$5.53 per barrel of crude oil exported to the U.S. This loss is equivalent to 40% of the average price differential in our sample. The Canadian oil and gas industry lost 2.74%–5.34% of its sales revenue and 68.72%–101.61% of the royalty payments because of insufficient pipeline capacity in 2015 and 2016. However, western Canadian oil refiners and consumers benefit from the depressed crude oil prices due to pipeline capacity constraints. Because changes in crude oil prices only partially pass through to changes in refined product prices,

a Department of Economics, University of Calgary, Calgary, AB T2N 1N4, Canada. E-mail: wdwalls@ucalgary.ca.

b Corresponding author, School of Business, Shanghai University of International Business and Economics, 1900 Wenxiang Rd, Shanghai, 201620, China, E-mail: xzhengsh@gmail.com.

western Canadian refiners increase profitability by \$0.1 per gallon of refined product they produce. Gasoline and diesel consumers benefited from the lower prices of refined products about \$0.2 per gallon. The total gains captured by local refiners and consumers are very small in comparison to the large losses of the upstream sector.

Our paper is particularly useful in understanding the consequence of recent policy changes in western Canadian oil industry. Alberta Premier Rachel Notley in December 2018 announced a temporary 8.7% oil production cut, a decrease of 325,000 barrels per day in the production of raw crude oil and bitumen starting January 2019.

On entry cost dynamics in Australia's National Electricity Market

Paul Simshauser^{ab} and Joel Gilmore^c

Historically spot prices in Australia's National Electricity Market (NEM) have exhibited considerable volatility within, and across, reporting periods. Along with short run variations associated with weather and anthropogenic patterns, medium-run supply-imbalances drive this volatility. Over the long run, given aggregate demand growth, or more relevantly in the current environment with flat final demand – the exit of aging coal plant “*at-scale*” – average spot prices will gravitate towards the cost of the relevant new entrant technology (or technology set). That is, higher prices on average, or during certain periods, will create incentives for targeted new entrant plant which in turn has the effect of capping longer-dated average spot price expectations at the estimated cost of the relevant new entrant technologies.

Over time, prices on average or during certain periods also regulate the plant mix as defined by the rich blend of fixed and variable costs associated with various generating technologies (i.e. base, intermediate, peak, variable renewables). Security-constrained power system simulation models reinforce this view. At their core, such models are based on equilibrium analysis.

Of course, in practice energy markets are frequently off-equilibrium. Near-term spot and forward contract prices can and do fall well below, or substantially exceed the relevant entry cost benchmarks and sometimes for extended periods due to transient structural imbalances within the plant stock. That structural imbalances exist in the first place means the cure to rising prices is not always more base plant. Understanding these principles is quite essential to understanding the fundamentals of power system planning, likely investment commitments and the long run marginal cost of power generation. Central to the task of power system modelling and investment analysis is therefore defining the equilibrium price of power, and for expediency we refer to as the new entrant cost.

In this article, we trace entry cost benchmarks for new generation plant and their relationship to spot price outcomes in Australia's National Electricity Market over the 20-year period to 2018. Over this period of time, the new entrant benchmark has shifted through a number of distinct phases involving transitions from coal, to natural gas, and more recently to “variable renewables plus firming”.

a Professor of Economics, Griffith Business School, Griffith University. Research Associate, Energy Policy Research Group, University of Cambridge.

b Head of Corporate Development at Infigen Energy.

c Head of Regulatory Affairs at Infigen Energy.

In the current phase, “variable renewables” means wind or solar PV, and “firming” is *notionally* (or shadow-) priced at the carrying cost of an Open Cycle Gas Turbine plant – either physically or financially through derivative instruments. Over the medium term this benchmark appears sound enough.

But over the long run, important implicit assumptions underpinning this particular (and notional) new entrant technology set may not hold if low marginal running cost coal plant continue to exit *at-scale*. The reason for this is that this benchmark relies critically on the gains from exchange in the NEM mandatory gross pool, efficiently utilising spare capacity. As aging coal plant exit, gains from exchange may gradually diminish with ‘notional firming’ increasingly and necessarily being met by physical firming. At this point, the benchmark must once again move to a new technology set.