

Griffith, Rachel; O'Connell, Martin; Smith, Kate

Working Paper

Price floors and externality correction

IFS Working Paper, No. W20/37

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Griffith, Rachel; O'Connell, Martin; Smith, Kate (2020) : Price floors and externality correction, IFS Working Paper, No. W20/37, Institute for Fiscal Studies (IFS), London,
<https://doi.org/10.1920/wp.ifs.2020.3720>

This Version is available at:

<https://hdl.handle.net/10419/242896>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Rachel Griffith
Martin O'Connell
Kate Smith

20/37

Working paper

Price floors and externality correction



erc

European Research Council



Economic
and Social
Research Council

Price floors and externality correction

Rachel Griffith, Martin O’Connell and Kate Smith*

November, 2020

Abstract

We study the introduction of a price floor for alcohol that is aimed at correcting for negative consumption externalities. Policy effectiveness depends on whether the measure achieves large reductions in the most socially costly consumption. We exploit a natural experiment to show the policy raised prices of cheap products favored by heavy consumers, and achieved large demand reductions among this group. We use pre-reform data to estimate a model of consumer demand that is able to match these patterns, and use this to compare the welfare performance of a price floor with the counterfactual introduction of an ethanol tax. We show that if the marginal external cost of drinking is at least moderately higher for heavy drinkers, then a price floor is better targeted at the most socially costly consumption and therefore achieves larger welfare gains than an ethanol tax. Although the price floor leads to a larger fraction of the consumer burden falling on those with low incomes compared with the tax reform, it leads to a consumer burden that is smaller for all income groups.

Keywords: externality, corrective taxes, alcohol, price floors

JEL classification: D12, D62, H21, H23

Acknowledgements: The authors would like to gratefully acknowledge financial support from the European Research Council (ERC) under grant number ERC-2015-AdG-694822 and from the Economic and Social Research Council (ESRC) under the Centre for the Microeconomic Analysis of Public Policy (CPP), grant number ES/M010147/1. Data supplied by Kantar FMCG Purchase Panel. The use of Kantar FMCG Purchase Panel data in this work does not imply the endorsement of Kantar FMCG Purchase Panel in relation to the interpretation or analysis of the data. All errors and omissions remain the responsibility of the authors.

*Griffith is at the Institute for Fiscal Studies and University of Manchester, O’Connell is at the Institute for Fiscal Studies and Smith is at the Institute for Fiscal Studies and University College London. **Correspondence:** rgriffith@ifs.org.uk, martin.o@ifs.org.uk and kate.smith@ifs.org.uk.

“The consumption of ardent spirits, particularly, no doubt very much on account of their cheapness, is carried to an extreme which is truly to be regretted, as well in regard to the health and morals as to the economy of the community.”

– Alexander Hamilton, First Report on the Public Credit (1790)

1 Introduction

Negative externalities from consumption are common, ranging from the social and health costs of drinking, smoking or drug use, to the environmental damage caused by fossil fuel use. In a simple textbook setting, a Pigouvian tax levied on the source of an externality can achieve the first-best allocation. However, the assumptions underpinning this result often break down in practice – importantly, in real world settings, externalities are often concentrated in a relatively small number of consumers. Price floors have been advocated as an effective policy to tackle problematic drinking (World Health Organization (2017)). Price floors are common in labor markets to redistribute to workers and agricultural markets to protect small upstream producers from monopsony power, but are banned in most other settings to protect consumers. Price floors can lower socially costly consumption by raising prices, but, unlike higher taxes, create windfall profits for firms instead of raising tax revenue, so have not been favoured by economists.

Our contribution in this paper is to study the impact of a price floor for alcohol on prices and quantities, and to compare its welfare effects with those of a tax levied on ethanol, the source of alcohol-related externalities. We exploit a natural experiment to show that the introduction of a price floor raised prices of cheap products favored by heavy consumers, and achieved large demand reductions among this group. We use pre-reform data to estimate a model of consumer demand that is able to match these patterns, and use this to compare the welfare performance of a price floor with the counterfactual introduction of an ethanol tax. When the marginal externality from drinking is roughly constant across drinkers the ethanol tax leads to larger welfare gains than the price floor, but when the marginal externality of heavy drinkers is moderately higher than for light drinkers the price floor does better than the tax. Even if the policymaker places zero value on firm profits, the superior targeting properties of the price floor outweigh the large transfer of surplus to the alcohol industry.

We begin by outlining a simple framework to compare the welfare impact of policies aimed at correcting for externalities. We compare a marginal increase in a

tax levied on the source of an externality with a marginal increase in a price floor in order to illustrate the intuition behind the relative performance of the two policies. The effect on welfare comprises gains from a reduction in externalities and losses from distorting consumption patterns. If the welfare criterion excludes firm profits, an increase in the price floor also entails a costly transfer of surplus to firms. On the other hand, if welfare is defined inclusive of firm profits this transfer cost does not apply. Crucial in determining the relative effectiveness of the two measures is whether they achieve reductions in the most socially costly consumption. If those consumers that create large *marginal* externalities are particularly likely to purchase products that are cheap, are relatively willing to switch away from these products, and are relatively disinclined to switch to more expensive alternatives, this will act to make the price floor relatively well targeted at externalities.

We study corrective policy in the alcohol market. The external costs of drinking are well established and substantial¹ and there is evidence of considerable variation in how costly drinking is across different people, with heavy drinkers responsible for a large share of externalities.² Specific alcohol taxes have long been used to tackle these externalities, but there is growing interest in using alternative price based policies. In 2018 Scotland introduced a price floor for alcohol, and Ireland has legislated for a similar policy. These price floors prohibit the sale of alcohol below a certain price per “unit of alcohol” (equivalent to 10ml of ethanol), and are explicitly aimed at tackling problem drinking. A number of Canadian provinces also have a system of minimum alcohol prices, though these are motivated as a way to limit the competition that state owned retailers face from private competitors.³

We exploit this natural experiment that entailed the introduction of an alcohol price floor in Scotland, but not in other parts of the United Kingdom.⁴ We use longitudinal micro data on the alcohol purchases from supermarkets and liquor stores of over 30,000 Scottish and English households before and after the reform. Figure 1.1 shows that average price and purchases moved similarly in Scotland and England prior to the reform, and that average prices rose sharply and quantity of alcohol purchased declined in Scotland post reform. Our difference-in-differences

¹For example, the Centers for Disease Control and Prevention (2016) estimate excessive alcohol consumption cost the US \$249 billion in 2010.

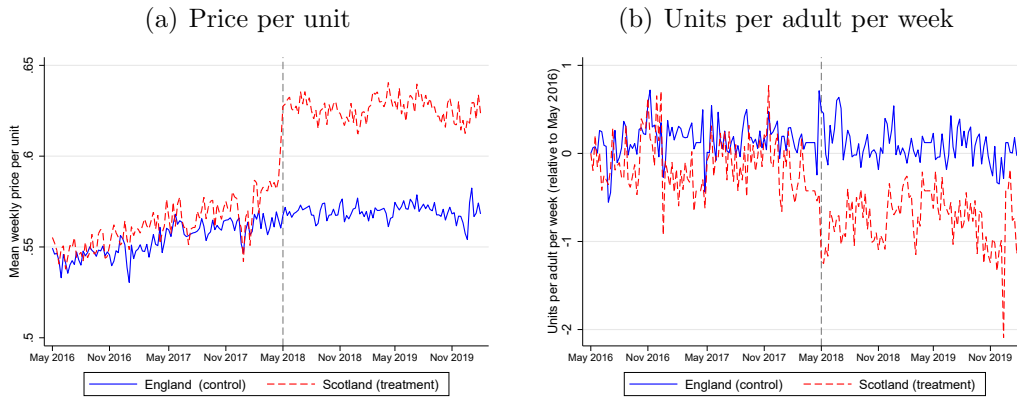
²See, for example, Centers for Disease Control and Prevention (2016).

³Although not the explicit aim of the policy, a number of studies have assessed their public health implications, finding a link between minimum pricing and lower alcohol consumption (Stockwell et al. (2012), Stockwell et al. (2012)), and an associated reduction in alcohol-related crime (Stockwell et al. (2017)) and morbidity (Zhao and Stockwell (2017)).

⁴The United Kingdom is made up of four “nations” – England, Scotland, Wales and Northern Ireland. Some policies are common across the UK, while other policies are devolved. Importantly for our application, taxes on alcohol are common across the nations of the UK.

estimates, which control for time and household fixed effects and possible contamination in the control group through cross-border shopping, show that the policy led to a 5% rise in average price per unit of alcohol across products, with some products experiencing price rises of more than 100% while the prices of others were unaffected. This shift in the price distribution led to an 11% fall in the average quantity of alcohol purchased. These price rises led to larger quantity reductions among heavy drinkers, in part because they obtained a greater share of their alcohol from products previously priced below the floor.

Figure 1.1: *Average alcohol prices and purchases in England and Scotland, 2016-20*



Notes: The left-hand panel shows the mean price paid per unit of alcohol, and the right-hand panel the mean units purchased per adult per week across households relative to the first week in May 2016, in England and Scotland in each week from May 2016 to January 2020. We remove country-specific week effects from all series. The dashed vertical line shows the introduction of the Scottish price floor on 1 May 2018. 1 unit of alcohol = 10ml of ethanol. For details of the data see Section 3.2.

We estimate a model of demand for alcohol products in order to compare the welfare impact of the price floor with an ethanol tax. Our motivation for using this approach instead of a sufficient statistics one (e.g., Chetty (2009)) is that we do not observe clear quasi-experimental variation that would allow us to identify the impact of a counterfactual, Pigouvian-style, ethanol tax. Estimating demand requires making a number of identification and functional-form assumptions. We estimate demand using data from an earlier time period, and we show that our demand model does a good job of predicting the effects of the price floor that we estimate under weaker assumptions using the quasi-experimental variation in the difference-in-differences analysis (see Angrist and Pischke (2010)).

We build on the discrete choice demand model in Griffith, O’Connell, and Smith (2019). Our model embeds the consumer’s decision over whether to buy alcohol, what type to buy and how much. We allow for preference heterogeneity across light, moderate and heavy drinkers (defined using pre-sample data) and household income.

This enables us to capture differences in demand responses across consumers who are likely to create very different levels of externalities and to analyze the distributional effects of policy. We estimate that the heaviest drinking households have higher own-price elasticities for individual alcohol varieties, but are also the most willing to switch between varieties when the price of one changes (and less likely to switch away from alcohol). As a result, their total demand for alcohol is substantially less elastic than the lightest drinkers. We simulate the effect of the Scottish price floor using our model, and find that it does a good job of predicting the average effect on quantity purchased, as well as variation across light and heavy drinkers and by different alcohol types when compared with our difference-in-differences results.

We use the model to compare the welfare impact of the price floor with that of an ethanol tax that achieves the same aggregate reduction in alcohol as the price floor. We calibrate the function mapping ethanol consumption into externalities, and, holding fixed total external costs, vary the function's convexity, which allows us to isolate how variation in the relative size of marginal externalities between heavy and light drinkers impacts the relative performance of the two policies. When the externality associated with each unit of alcohol consumed is constant, the ethanol tax out-performs the price floor – overall it increases welfare by £198 million per year, while the price floor leads to a fall of £225 million. However, when there is even a moderate degree of convexity in the externality function, the price floor leads to larger welfare gains than the ethanol tax. For instance, when heavy drinkers – who consume around half of all ethanol – account for 85% of the external costs of drinking, the price floor raises welfare by £470 million per year more than the ethanol tax. This is driven by the price floor being much better targeted at heavy drinkers, and thus leading to a larger reduction in external costs, which more than offsets the lower tax revenue.

Our results are related to our previous work, Griffith, O'Connell, and Smith (2019), which also considers the effectiveness of policy when there are heterogeneous externalities. In this setting a single tax rate cannot achieve the first best and, as highlighted by Diamond (1973), must trade-off price rises that are too low for the most socially costly consumers and too high for low cost consumers. In that paper we show how varying tax rates across product types can create efficiency gains when product level demands are correlated with marginal externalities. Relatedly, there is an environmental literature that focuses on the challenge of designing policy when it is difficult to directly target the source of the externality, and that compares the efficacy of targeting different product features.⁵ Here we study how the price at

⁵See, for instance, Grigolon et al. (2018) and Jacobsen et al. (2020).

which a product that supplies an externality generating characteristic can be used to “tag” socially costly consumption. A benefit of targeting this via a price floor is that it sidesteps the problems associated with “line drawing” (Gillitzer et al. (2015)) that may hamper the implementation of increasingly complex systems of excise taxes. Variation in externalities is not unique to alcohol consumption; for example, daily cannabis use substantially increases the risk of psychosis and other mental impairments (World Health Organization (2016)). Our findings suggest that price floors could be a useful tool for governments as they design externality correcting policy.

Another strand of the literature on externality correcting policy considers how a planner’s preference over redistribution may impact policy; for example, see Allcott et al. (2019). We use a welfare criterion that includes a money metric measure of consumer surplus; our welfare comparisons are therefore not impacted by redistribution across consumers. It is nonetheless important that we assess the redistributive effects of reforms, to inform the nature of any compensatory adjustments that policymakers may wish to make through other parts of the tax system. We allow preferences to vary by household income, which enables us to compare the change in consumer surplus under each policy across the income distribution. A common criticism of the alcohol price floor is that it disproportionately burdens the poorest consumers. We show that, although low income households bear a greater *proportion* of the reduction in consumer surplus under the price floor than the ethanol tax, the fall in aggregate consumer surplus is much lower under the price floor because it is better targeted. Thus, the fall in consumer surplus for low income households is less under the price floor than under the ethanol tax.

A number of recent papers study the effects of public policy in alcohol markets, with a particular focus on the role of firms.⁶ Unlike in many US states, where alcohol retailing is monopolized by a state owned retailer, in the UK, alcohol sales are dominated by the large supermarkets. We provide evidence that, in this market, taxes are passed through one-for-one to prices, and under the price floor, products formerly priced below the floor increase to the floor and the prices of other products are largely unaffected. We use a welfare criteria that does not include firm profits. This means windfall profits conferred on firms through a price floor are not valued

⁶Seim and Waldfogel (2013) show that in Pennsylvania replacing the state monopoly with a private market would lead to a much expanded supply of liquor stores, while Miravete et al. (2018) and Miravete et al. (2020) consider the revenue maximizing mark-up for the state retailer and show revenues could be considerably higher if the mark-up varied across products. Conlon and Rao (2019) consider the effect of post-and-hold regulations in Connecticut and show they help wholesalers sustain prices closer to the collusive level, and removing the regulation and adjusting tax rates could result in gains in tax revenue and consumer surplus.

by the policymaker. We discuss the implications of including them, when comparing the price floor and ethanol tax, showing that it reinforces the better performance of the former over the latter (under a convex externality function).

The rest of the paper is structured as follows. In the next section, we set out a simple framework for comparing the welfare effects of policies aimed at externality correction. In Section 3 we describe the Scottish price floor and use a difference-in-differences approach to estimate its impact on prices and quantities. In Section 4 we set out and present estimates of a model of alcohol demand, and show it can match the effects presented in Section 3. In Section 5 we use the model to compare the welfare effects of the price floor with a counterfactual ethanol tax. A final section summarizes and several appendices provide additional detail.

2 Motivating example

We begin by outlining the marginal welfare effects of policies aiming to correct for externalities, in order to provide some intuition of what drives our results below. The discussion in this section highlights the main forces that determine the effectiveness of a price floor, relative to taxation, and indicates the key objects that we need to empirically estimate to assess the relative performance of different policies.

Setup. A population of heterogeneous consumers, indexed i , choose between a set of externality generating products, $j = \{1, \dots, J\}$ (for example, different alcohol products). For ease of exposition, we assume each consumer’s utility is quasi-linear and is given by $v_i(\mathbf{p}) + y_i$, where $v_i(\cdot)$ is the monetary surplus consumer i gets from participating in the market, $\mathbf{p} = (p_1, \dots, p_J)$ are prices and y_i is consumer income. Denote consumer i ’s demand for product j by q_{ij} , their demand for all products in the market $Q_i = \sum_j q_{ij}$ and total consumer demand for product j by $Q_j = \sum_i q_{ij}$.

Consumption of the products $j = \{1, \dots, J\}$ generates externalities. We normalize quantities and prices so they are expressed per unit of the externality generating characteristic. In the case of alcohol quantities are in terms of “units of alcohol”, and prices are per unit (a unit of alcohol equals to 10ml of ethanol, and is the standard unit of measurement in Europe). We denote the total externality associated with consumer i by $\Psi_i(Q_i)$. In many cases $\Psi_i(\cdot)$ will be a weakly increasing, weakly convex function that goes through the origin.

Suppose the government has access to two policy instruments, (i) a linear tax rate, τ , and (ii) a price floor, \underline{p} , where each instrument applies directly to the source of the externality. In the case of alcohol the tax rate is levied, and the

price floor defined, per alcohol unit. We assume that policy is translated one-for-one into prices. In particular, letting c_j denote the price of product j in the absence of government policy, the consumer price is $p_j = \max\{\underline{p}, c_j + \tau\}$. We offer direct evidence below that this assumption holds in the market we study. A market structure that would generate this conduct is perfect competition and constant returns to scale technology (in which case c_j is product j 's marginal cost). Whether policy is translated one-for-one to prices under imperfect competition is an empirical question.⁷

In comparing the two policies, we base our welfare criterion on the sum of consumer surplus and tax revenue minus externalities:

$$W = \sum_i v_i(\mathbf{p}) + \sum_j \tau Q_j - \sum_i \Psi_i(Q_i).$$

An alternative is to also include profits in the criterion, which would accrue under a price floor, even under a perfectly competitive market structure. We comment below on how this would affect the comparison of the two policy instruments.

Marginal welfare effects. In this simple example we compare the impact on welfare that results from a marginal rise in the tax rate with a marginal rise in the price floor. Let \mathcal{B} denote the set of products for which the price floor binds.⁸

Consider first a marginal increase in the tax rate. Note that with a price floor in place, a marginal tax increase will raise only the price of those products priced above the floor. For the set of products in \mathcal{B} the tax rise will transfer part of the gross margin on the product, $\underline{p} - c_j$, from firms to the government. The impact on welfare of the marginal tax rise can be written:

$$\frac{dW}{d\tau} = - \sum_i \frac{d\Psi_i}{dQ_i} \frac{dQ_i}{d\tau} + \sum_j \tau \frac{dQ_j}{d\tau} + \sum_{j \in \mathcal{B}} Q_j.$$

The first term, $-\sum_i \frac{d\Psi_i}{dQ_i} \frac{dQ_i}{d\tau}$ is a “corrective” term. It captures the increase in welfare associated with the tax rise lowering externalities. $\sum_j \tau \frac{dQ_j}{d\tau}$ is the “revenue leakage” term, which is given by the loss in government revenue resulting from consumers modifying their behavior in response to the tax rise. $\sum_{j \in \mathcal{B}} Q_j$ is a “transfer”

⁷In the highly concentrated sugar sweetened beverage market, there is evidence of full pass-through of taxes (see Seiler et al. (2020), Dubois et al. (2020)).

⁸For ease of exposition we assume that a marginal increase in either policy does not change the set of products subject to the price floor. This means a marginal increase in the price floor raises the price of all products in the set \mathcal{B} , but does not cause any additional products to be subject to the price floor. This would be the case if, at the current price floor, there exist no products with price $p \in \{\underline{p}, \underline{p} + \epsilon\}$. This assumption is not crucial when considering marginal policy changes, but significantly simplifies the algebra.

term. It reflects a welfare gain from the transfer from firms to government which comes from products subject to the price floor.⁹ If no binding price floor is in place, the transfer term disappears, and the tax rise will increase the price of all products.

A marginal increase in the price floor raises the price of those products subject to the floor (i.e., cheap sources of externality), but not those priced above it. The price rises translate into increases in the margin on these products, transferring surplus to firms. The impact of a marginal increase in the price floor on welfare, like that for a tax rise, comprises the sum of corrective, revenue leakage and transfer terms and is given by:

$$\frac{dW}{d\underline{p}} = - \sum_i \frac{d\Psi_i}{dQ_i} \frac{dQ_i}{d\underline{p}} + \sum_j \tau \frac{dQ_j}{d\underline{p}} - \sum_{j \in \mathcal{B}} Q_j.$$

The corrective and revenue leakage terms for the price floor and tax rise differ because each policy impacts the prices of different products: the increase in the price floor raises the price of relatively cheap products, the tax rise increases the price of all products not subject to a binding floor. The transfer term takes the same form except with its sign reversed; the transfer to firms conferred by an increase in the price floor lowers welfare.¹⁰

When welfare is based on consumer surplus and tax revenue net of externalities, a clear disadvantage of a price floor relative to tax is that it transfers surplus from consumers to firms (in contrast to tax, which transfers it from consumers to government).¹¹ However, this disadvantage may be outweighed if the price floor is a more effective corrective instrument. This is more likely to be the case the more highly correlated are the marginal externalities created by consumers, $d\Psi_i/dQ_i$, and the size of reductions in their total consumption the price floor rise achieves, $dQ_i/d\underline{p}$. Hence, the more that the most socially costly consumers (i) get a large share of their consumption from cheap sources of externality (those affected by the price floor), (ii) have a large own-price elasticity for these products, and (iii) have

⁹Note if we use a welfare criterion that places weight on firms' profits this term is replaced with one that reflects the loss in profits to firms accruing from consumers adjusting their demands. For instance, under perfect competition, the term $\sum_{j \in \mathcal{B}} Q_j$ would be replaced with $\sum_{j \in \mathcal{B}} (w_j - c_j) \frac{dQ_j}{d\underline{p}}$, where $w_j = \max\{\underline{p} - \tau, c_j\}$ is the tax exclusive price.

¹⁰However, as with the tax rise, the transfer term is replaced by an efficiency term when firm profits are included in the welfare criterion; under perfect competition the term $-\sum_{j \in \mathcal{B}} Q_j$ is replaced with $\sum_{j \in \mathcal{B}} (w_j - c_j) \frac{dQ_j}{d\underline{p}}$.

¹¹This disadvantage disappears if firm profits are included in the welfare function.

a low cross-price elasticity for products priced above the floor, the better the price floor will be at correcting for externalities.¹²

Summary. This discussion guides our empirical analysis of the performance of a price floor in the alcohol market. First, we use quasi-experimental variation that arises from the introduction of the policy in Scotland but not in England to estimate its impact on prices and quantities. We show that the effects vary across light and heavy drinkers. This is important because marginal externalities are likely to be higher, on average, for heavier drinkers. As discussed above, this differential impact across the drinking distribution is a key determinant of the impact of a price floor relative to taxation. Second, we estimate a model of consumer demand, which we need to consider the counterfactual introduction of an ethanol tax. We show that the out-of-sample predictions of the impact of the introduction of a price floor using our demand model are consistent with those estimated using the quasi-experimental variation. We then use the model to compare the welfare implications of the introduction of a price floor with the introduction of an ethanol tax.

3 Estimating the effect of an alcohol price floor on price and quantities

3.1 Policy context

A price floor for alcohol – known as a minimum unit price – came into effect in Scotland on May 1, 2018; this policy did not apply in England. The policy made it illegal to sell alcohol products priced below a floor equal to £0.50 per unit of alcohol. A unit of alcohol is 10ml of ethanol and is the standard metric in the UK and many European countries.¹³ The Scottish Parliament first legislated for the policy in 2012, but its implementation was delayed due to a legal challenge.¹⁴

¹²To see this, note we can re-write $\frac{dQ_i}{dp} = \frac{Q_i}{p} (\eta_{\mathcal{B}} s_{\mathcal{B}} + \eta_{\mathcal{B}^C} s_{\mathcal{B}^C})$ where $\eta_{\mathcal{B}} = \frac{d \sum_{j \in \mathcal{B}} q_{ij}}{dp} \frac{p}{\sum_{j \in \mathcal{B}} q_{ij}}$ is the own price elasticity of products in \mathcal{B} , $\eta_{\mathcal{B}^C} = \frac{d \sum_{j \notin \mathcal{B}} q_{ij}}{dp} \frac{p}{\sum_{j \notin \mathcal{B}} q_{ij}}$ is the cross price elasticity of products not in \mathcal{B} with respect to a price change of those in \mathcal{B} and $s_{j \in \mathcal{X}} = \frac{\sum_{j \in \mathcal{X}} q_{ij}}{Q_i}$ is the quantity share of products belonging to the set $\mathcal{X} = \{\mathcal{B}, \mathcal{B}^C\}$.

¹³The convention for measuring ethanol volume in the US is a “standard drink”, which is 17.7 ml of ethanol.

¹⁴The Scottish Whisky Association argued the measure would encourage trade discrimination. The Scottish Government ultimately prevailed as they successfully overcame the hurdle placed by the Court of Justice of the European Union that a “practice such as that adopted in Scotland is not justified where it is possible for health to be protected equally effectively by less restrictive tax measures” (Court of Justice of the European Union (2015)).

The policy is motivated as a means of tackling externalities from alcohol consumption: the devolved Scottish Government, “*wants to target the price of drinks that are cheap and strong,*” as these are, “*the alcoholic drinks that tend to be drunk by people who are at more risk of harm due to drinking*” (Scottish Government (2018)). There is considerable evidence on the social costs of excess alcohol consumption (World Health Organization (2014)), which include public healthcare costs as well as the effects of drink driving, domestic violence and other crime. Much of this evidence suggests that there are “threshold effects” of drinking: the risk is minimal at low levels of alcohol consumption, but rises sharply when consumption exceeds low levels. For instance, there is evidence of a threshold effect in the relationship between alcohol consumption and the risks of developing tuberculosis (Lönnroth et al. (2008)), and liver cirrhosis (Rehm et al. (2010)). Convexity in the relationship between alcohol consumption and harms is not limited to disease (and associated public health costs): for instance, harmful and hazardous levels of alcohol consumption have been shown to significantly raise the risk of perpetrating domestic violence (World Health Organization (2006)). This evidence is reflected in government guidelines on drinking, e.g., in the UK, people are advised not to consume more than 14 units per adult per week, which, while not regarded as “safe”, is termed “low risk” (National Health Service (2018)).

Alcohol sold in Scotland is also subject to taxes that are set by the UK government (and are therefore the same in Scotland and England); see Appendix C.1 for a description. Products are taxed both under a system of alcohol duties and a general consumption (value added) tax (VAT). The alcohol duties include a volumetric tax (i.e., per liter) for wines and ciders, and a specific tax on ethanol for beer and spirits. VAT is levied at a rate of 20% on the duty inclusive price. The system of alcohol duties is broadly similar to those in European countries. It also shares key features with US alcohol duties, such as variation across alcohol types with higher rates on spirits, although US tax rates are levied at a lower rate. In Section 5 we show that our conclusions about the relative performance of the policies we consider are robust to varying the baseline level of taxation.

3.2 Data

We use data from the Kantar FMCG Purchase Panel, which is a household level scanner dataset collected by the market research firm Kantar UK. A representative sample of UK households record all grocery purchases they make and bring into the home. This includes purchases from supermarkets, convenience stores and liquor stores. Table A.1 in the Appendix shows that the sample is similar along key

demographics with the nationally representative consumer spending survey, the Living Costs and Food Survey (LCFS). Households record the products (at the barcode or UPC level) that they buy, along with transaction level prices; Kantar also collect information on product and household characteristics. The data are longitudinal, with households typically present in the data for several years.

We evaluate the impact of the Scottish price floor on prices and purchases using data from May 2016 to January 2020, covering 24 months prior to and 20 months following the introduction of the policy. The data cover over 2.6 million alcohol purchases made by 32,480 households living in Scotland and England. We observe households for an average of 115 weeks over this period. Table A.1 shows that, conditional on making an alcohol purchase over a two-week period, households in the Kantar data report buying a similar quantity of units per adult per week to those in the Living Costs and Food Survey.¹⁵ We use similar data from an earlier period to estimate consumer demand, see Section 4.1.

An important limitation of these data is that they do not include information on alcohol bought for consumption out of the home, i.e., in restaurants and bars. Alcohol bought for at-home consumption accounts for around three-quarters of alcohol units consumed in the UK (data from the LCFS). The prices of alcohol purchased in restaurants and bars will not directly be affected by the price floor, as it is priced well above this level, but consumers may respond to the introduction of the price floor by switching from at-home consumption to consumption out of home. However, a recent study by the Scottish Government (NHS Health Scotland (2019)) using aggregate alcohol sales data finds that there was no difference between Scotland and England in the change in the quantity of alcohol purchased in restaurants and bars between 2017 and 2018. We further discuss the implications of omitting this source of alcohol for our conclusions in Section 5.

We estimate the effect of the price floor in Scotland using a difference-in-differences approach, with England as a control group. A potential issue is that cross-border shopping could contaminate our estimates. Our data are recorded at the household level, so if a Scottish household buys alcohol in bordering English counties, then we will incorrectly consider these transactions to be subject to the price floor. There is some evidence of this, which we discuss in Appendix A.2. To avoid this biasing our results we exclude households living within 50km of the bor-

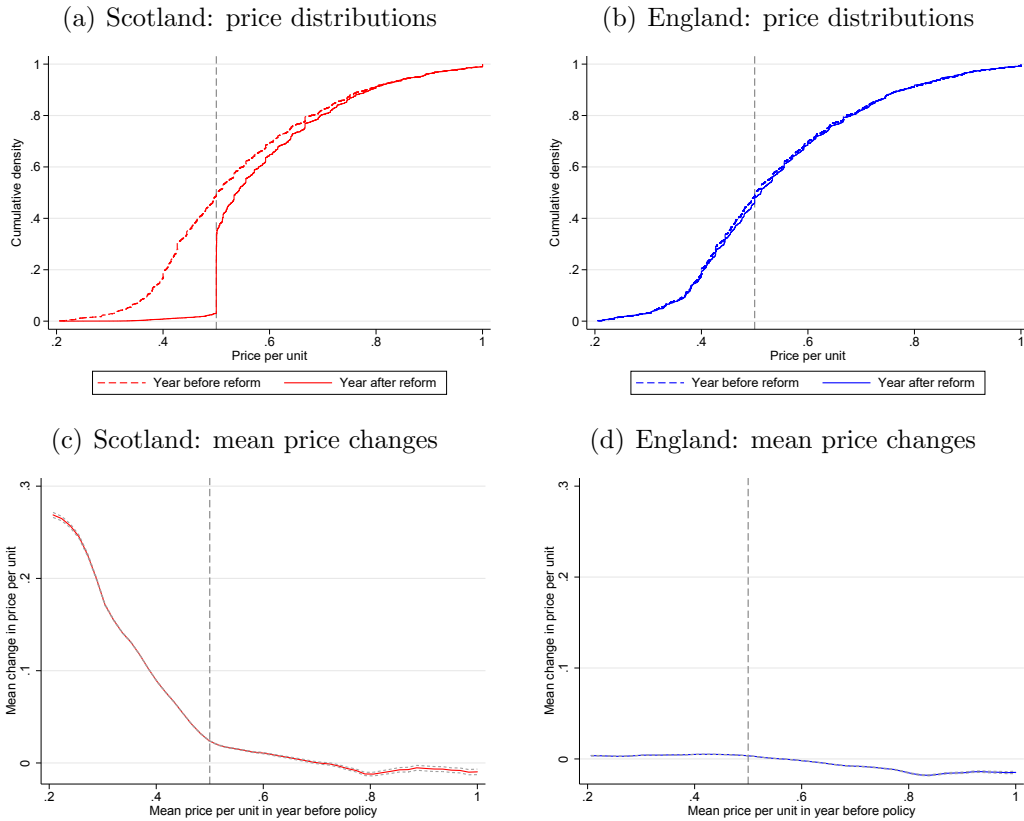
¹⁵We condition on buying alcohol in order to make the measurement comparable across the two datasets: the LCFS records household spending using a two-week diary, which means that we cannot compute long-run measures of household alcohol purchases in the LCFS. We also compare the share of households who record never buying alcohol in the Kantar data to share who report not drinking in the past twelve months in the Health Survey for England; both record approximately 15%.

der between Scotland and England. This area is sparsely populated, which means we drop only 1% of households in our sample. This leaves us with 2,856 Scottish households in our treatment group and 29,306 English households in our control group.

3.3 Impact on prices

Let ρ_{jt} denote the price paid for product j on transaction t , and z_j the number of units of alcohol in product j . There are 12,463 products (or barcodes) in our data. We use $p_{jt} = \rho_{jt}/z_j$ to denote the price per unit of alcohol. Figure 3.1 summarizes how the distribution of prices changed in Scotland and England from the year before to the year after the introduction of the price floor.

Figure 3.1: *Effect on price distributions*



Notes: Panels (a) and (b) show the distributions of price paid per unit across transactions in the year before and the year after the introduction of the price floor in the Scotland and England, respectively. Panels (c) and (d) show, for the set of products that are recorded as purchased in the year before and after (which account for 80% of spending across the two years), the average change in price per unit, conditional on the product's average price in the year preceding the reform.

Panel (a) shows that just under 50% of transactions in Scotland were below the floor in the year before the reform; following the reform around 40% of transactions were exactly at the price floor. In comparison, panel (b) shows that in the year

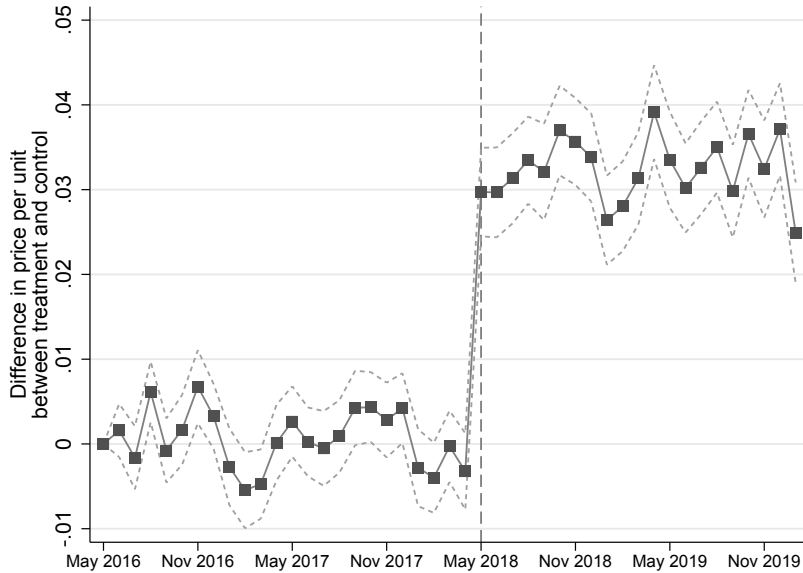
prior to the introduction of the policy, the transaction price distributions in Scotland and England were similar, and there was little change in prices in England over the same period. Panel (c) shows the mean price change in Scotland by the average price per unit in the year prior to the introduction of the policy. This shows the differential effect of the policy across the price distribution: some very cheap products experienced price increases in excess of 100%, while products that were priced above the floor pre-reform exhibit very little change in price. Panel (d) shows that we did not see these changes in England.

We use a difference-in-differences approach to assess the impact of the introduction of the price floor in Scotland, by comparing to price changes in England. We estimate:

$$p_{jt} = \sum_{m=2}^{45} (\beta_m \times \text{treat}_t + \gamma_m) + \zeta_j + \boldsymbol{\chi}_t + \epsilon_{jt} \quad (3.1)$$

where m indexes year-months, treat_t is a dummy variable equal to one if transaction t was made by a household that lives in Scotland, ζ_j denote product fixed effects, and $\boldsymbol{\chi}_t$ are controls for the weeks before Christmas, New Year and Easter, where we allow the effect of these holidays to differ between Scotland and England.¹⁶

Figure 3.2: *Difference in price per unit between Scotland and England*



Notes: The markers show the estimated $\hat{\beta}_m$ s from equation (3.1), estimated based on a sample of 2,835,499 transactions. The vertical dashed line indicates the month in which the price floor was introduced in Scotland. Price per unit is expressed in £.

¹⁶For instance, the tradition of celebrating “Hogmanay” – the last day of the year – in Scotland means considerably more alcohol per person is purchased there than in England.

Figure 3.2 plots the $\hat{\beta}_m$ s, which are the estimated within-product differences in price per unit between Scotland and England in the months before and after the introduction of the price floor. Prior to the introduction of the policy, the coefficients are not statistically different from zero, indicating that prices satisfy parallel pre-trends. Once the policy came into force, there was a sharp £0.03 /unit increase in mean price in Scotland relative to England. This corresponds to a mean price increase of roughly 5%, although, as shown in Figure 3.1 this mean effect consists of a mix of large price increases for cheap (per unit of alcohol) products and no changes for more expensive (per unit of alcohol) products.

3.4 Impact on purchases

The price floor led to significant increases in the price of previously cheap alcohol. We estimate the impact this had on the amount of alcohol purchased by households over this period. We define the number of alcohol units purchased per adult per week by household i in year-week w as:

$$Q_{iw} = \frac{1}{A_i} \sum_j z_j \left(\sum_{t \in T_{iw}} \eta_{jt} \right) \quad (3.2)$$

where z_j denotes the number of units in product j , η_{jt} the number of packs of product j bought on transaction t , T_{iw} the set of transactions made by household i in year-week w , and A_i the number of adults in household i . If a household records making any grocery purchases in week w , but does not buy any alcohol, we set Q_{iw} to zero.

Mean effect

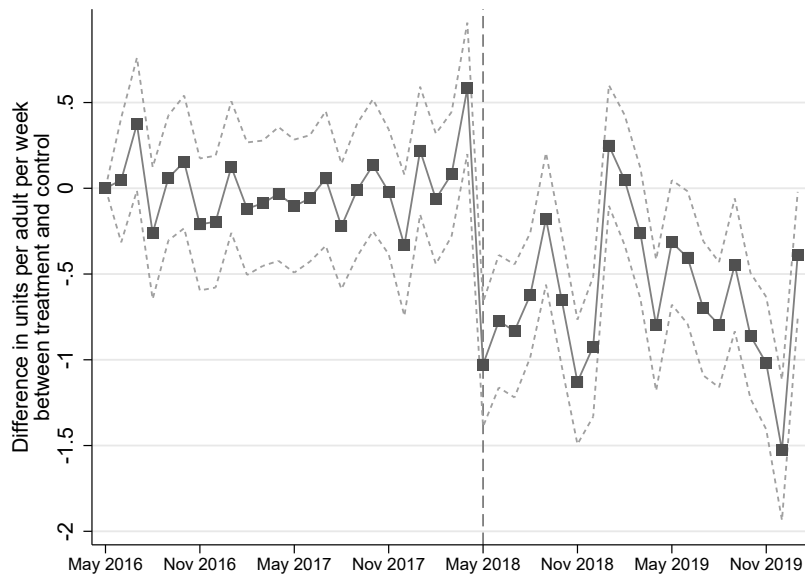
In Figure 1.1(b) we show the time series of mean units per adult per week in Scotland and England – the figure shows a clear decline in units purchased in Scotland following the introduction of the price floor. Here we estimate this difference over time within household by including household fixed effects. We estimate:

$$Q_{iw} = \sum_{m=2}^{45} (\beta_m \times \text{treat}_i + \gamma_m) + \mu_i + \chi_{iw} + \epsilon_{iw} \quad (3.3)$$

where treat_i is a dummy variable equal to one if household i lives in Scotland, μ_i denote a set of household fixed effects, and χ_{iw} denote, as above, controls for the weeks before Christmas, New Year and Easter, where we allow the effect of these holidays to differ between Scotland and England.

Figure 3.3 plots the $\hat{\beta}_m$ s, which are the estimated differences in the quantity of alcohol purchased (mean units per adult per week) between households in Scotland and England in the months before and after the introduction of the price floor. The coefficients are not statistically different from zero in the pre-period, indicating that units purchased evolved similarly in Scotland and England prior to the policy's introduction. There was an increase in the units purchased by Scottish households in April 2018; this is consistent with stockpiling behavior in anticipation of higher future prices. Following the introduction of the policy in May 2018, there was a drop in units purchased in Scotland relative to England. The month-to-month changes are somewhat noisy; however, overall there is a clear decline in Scotland relative to England from the period before to after the reform. The decline is 0.63 units per adult per week, which is highly statistically significant and represents an 11% reduction in average alcoholic units purchased (see Table A.2 in Appendix A.3).

Figure 3.3: *Difference in units per adult per week between Scotland and England*



Notes: The markers show the estimated $\hat{\beta}_\tau$ s from equation (3.3), estimated based on a sample of 4,020,485 household-year-weeks. The vertical dashed line indicates the month in which the price floor was introduced in Scotland.

Heterogeneity across drinker type

As discussed in Section 2, the efficiency gains achieved by the price floor depend on how changes in households' alcohol purchases are correlated with the external costs associated with their drinking. In Section 3.1, we describe some of the evidence that suggests that heavy drinkers have higher marginal externalities from drinking.

We use May 2016 – April 2017 to create a measure of long-run average purchasing behavior. We refer to this as the pre-sample period and estimate the effects of the price floor using the period May 2017 – January 2020. It is likely that the external costs of drinking are higher for individuals and households with higher long-run average alcohol purchases.

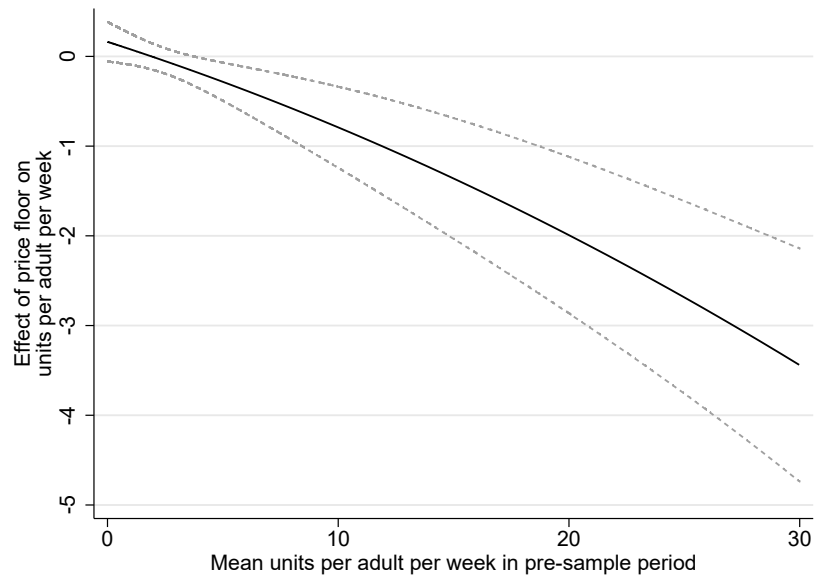
We allow the effect of the policy to be a quadratic function of the average number of units the household purchased in the pre-sample period. Let $\bar{Q}_i = \frac{1}{\mathcal{W}_i} \sum_{w=1}^{52} Q_{iw}$ denote the average number of units purchased per adult per week in the pre-sample period, May 2016 to April 2017, where \mathcal{W}_i denotes the number of weeks that the household recorded buying groceries in that year. We estimate:

$$Q_{iw} = (\beta_0 + \beta_1 \bar{Q}_i + \beta_2 \bar{Q}_i^2) \times \text{treat}_i \times \text{post}_w + \sum_{m=13}^{45} \gamma_m + \mu_i + \chi_{iw} + \epsilon_{it} \quad (3.4)$$

for May 2017 to January 2020 (i.e., excluding the period used to define \bar{Q}_i); post_w is a dummy equal to one if year-week w is after the implementation of the price floor; all other variables are defined as in equation (3.3).

We estimate (3.4) using a sample of 24,372 households that we observe in the pre-sample period; this sample of households is similar to the full sample, both in the distribution of alcohol units purchased and the average impact of the price floor. Figure 3.4 plots $\hat{\beta}_0 + \hat{\beta}_1 \bar{Q}_i + \hat{\beta}_2 \bar{Q}_i^2$, showing how the estimated effect of the price floor on units per adult per week varies with pre-sample purchases. The impact of the policy was largest for those households who consistently bought large amounts of alcohol in the pre-period, both in levels and proportionally. Those who bought, on average, five units per adult per week in 2016 experienced a reduction of 6%, compared with a fall of 12% for those who bought 30 units per adult per week. This suggests that the policy is effective at reducing the purchases of those whose purchases are highest, and whose consumption is more likely to create the largest externalities.

Figure 3.4: *Effect of the price floor across the distribution of drinkers*



Notes: The vertical axis plots $\hat{\beta}_0 + \hat{\beta}_1 \bar{Q}_i + \hat{\beta}_2 \bar{Q}_i^2$ (see equation (3.4)) estimated on a sample of 2,699,021 household-year-weeks. The horizontal axis plots mean units purchased per adult per week over May 2016 to April 2017, \bar{Q}_i . Dotted lines show 95% confidence intervals. See Table A.3 for the coefficient estimates.

Alcohol basket composition and substitution responses

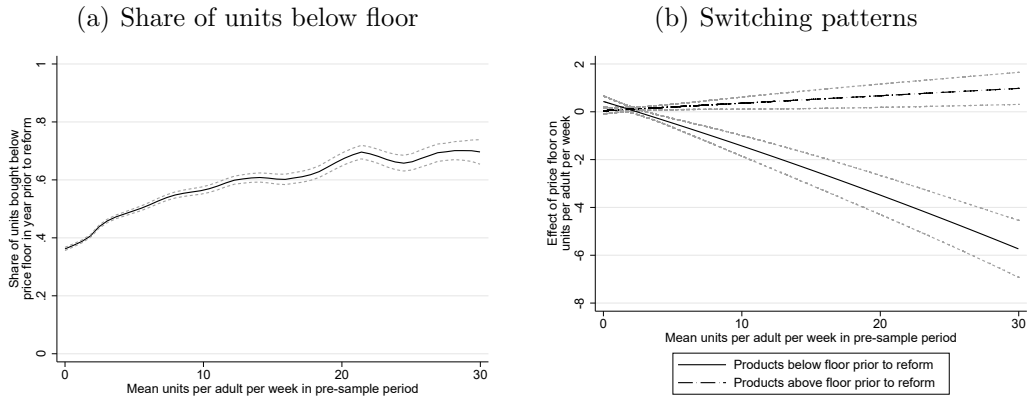
The effect of the price floor on the quantity of alcohol that households purchase is driven both by how much of their alcohol basket prior to the introduction of the policy was priced below the floor, and how they respond to the price increases.

Products that were priced below the floor prior to the introduction of the policy experienced an increase in their prices, both in absolute terms, and relative to other alcohol products. Those products that were priced above the floor see little absolute change in their price, and falls in their relative price. We estimate the difference-in-differences specification, equation (3.3) separately on (i) units from products previously priced below the floor and (ii) units from products previously priced above the floor.¹⁷ We find that there was a fall of 0.94 units per adult per week from products previously priced below the floor; this was partially offset by an increase of 0.26 units from products previously priced above the minimum (see Table A.2 in Appendix A.3).

¹⁷For each product we calculate its mean price per unit in the year prior to the introduction of the floor, and then we use this to divide products into those above and below the floor. For each household-year-week we sum units from each set of products.

Figure 3.5 shows how these patterns vary over the distribution of drinkers, which we measure as average units per adult per week in the pre-sample period (i.e., \bar{Q}_i). Panel (a) shows that in the year prior to the introduction of the price floor, heavier drinkers (i.e., those with high values of \bar{Q}_i) bought a larger share of their alcohol from products priced below the floor – rising from 40% for the lightest drinkers to around 70% for the heaviest drinkers. Panel (b) shows how units from products previously priced below and above the minimum price changed for lighter and heavier drinkers. Specifically it plots how the estimates of $\hat{\beta}_0 + \hat{\beta}_1 \bar{Q}_i + \hat{\beta}_2 \bar{Q}_i^2$ from equation (3.4), estimated separately for products previously priced below and priced above the price floor (Table A.3 in Appendix A.3 summarizes the coefficient estimates), vary with \bar{Q}_i . Households across the drinking distribution increased units from products priced above the minimum price pre-reform, and reduced units from products previously below the price floor. The relative slopes of the two lines indicate why the policy was relatively well targeted – heavier drinkers switched away from cheap products, but increased their purchases of more expensive products proportionately less than lighter drinkers.

Figure 3.5: *Switching across the distribution of drinkers*



Notes: The vertical axis on the left hand panel plots the share of units purchased below the price floor over the period May 2017 to April 2018 (in both England and Scotland). The vertical axis in the right hand panel plots $\hat{\beta}_0 + \hat{\beta}_1 \bar{Q}_i + \hat{\beta}_2 \bar{Q}_i^2$ based on equation (3.4), with two different dependent variables – units from products priced below the floor prior to the reform, and units from products priced above the floor prior to reform – estimated based on a sample of 2,699,021 household-year-weeks. In both panels the horizontal axis plots mean units purchased per adult per week over May 2016 to April 2017, \bar{Q}_i . In both panels the dotted lines show 95% confidence intervals. See Table A.3 for the coefficient estimates.

4 A model of consumer demand for alcohol

In order to compare the price floor with an ethanol tax, and to evaluate the welfare impacts, we estimate a model of household alcohol demand and use this to simulate counterfactual policy reforms. We use data prior to the introduction of the price floor, and show that the out-of-sample predictions of the demand model match well the impact of the price floor estimated using a difference-in-differences approach, described in Section 3. Our model and estimation approach follows closely our previous analysis in Griffith, O’Connell, and Smith (2019).

4.1 Demand specification

Data

We use data from the Kantar FMCG Purchase Panel (see Section 3.2) for January 2010 - December 2011. This is immediately before the Scottish Parliament legislated to introduce a price floor, and so our demand estimates will be unaffected by the policy. We estimate demand on a sample of 9,428 households that we observe buying alcohol in 2010 and 2011.¹⁸ We use the 2011 data to estimate demand and the 2010 data to measure households’ average long-run alcohol purchases; we refer to 2010 as the pre-sample period. In Appendix A we include more details about the data, and show that this sample is similar to the one used to estimate the effect of the price floor in Section 3.

Alcohol varieties

Price based alcohol policies, such as taxes and price floors, typically change the relative prices of different products, and these relative price changes are likely to lead consumers to change the composition of their alcohol baskets.

We follow the approach taken in our earlier work, Griffith, O’Connell, and Smith (2019), which we found to be a parsimonious way to capture consumer switching across different alcohol products. We specify a model of demand for varieties of alcohol, shown in Table 4.1. The 69 varieties vary by alcohol type (e.g., strong premium beer or budget whisky) and by size (e.g., 500ml, 1-2 litres or 2x700ml). For each variety we compute a price index that varies over months and reflects price movements in the underlying products (UPCs) that make up the variety. We use regional weights that are time-invariant, and reflect the share of transactions accounted for by each UPC in the alcohol variety.

¹⁸We observed households for an average of 40 weeks and a minimum of 20 weeks in each year.

Table 4.1: *Alcohol varieties*

(1)	(2)	(3)	(4)	(5)
Variety	Size	Mean quantity (L)	Mean price (£)	Mean price per unit (£)
<i>Beer</i>				
Premium beer; ABV < 5%	500ml	0.52	1.59	0.70
	1-2L	1.32	3.96	0.65
	2.5-8L	3.63	9.01	0.57
Premium beer; ABV ≥ 5%	500ml	0.52	1.82	0.61
	1-2L	1.35	4.19	0.56
	2.5-10L	3.59	9.92	0.50
Mid-range bottled beer	1-2L	1.43	3.50	0.52
	2.5-4L	3.01	6.71	0.47
	5-14L	6.58	12.79	0.43
Mid-range canned beer; ABV < 4.5%	2-5L	3.47	6.15	0.45
	7-10L	8.13	12.80	0.40
	15-25L	14.72	19.93	0.40
Mid-range canned beer; ABV ≥ 4.5%	1-3L	1.94	4.43	0.46
	4-6L	4.16	8.58	0.41
	8-20L	8.94	15.96	0.39
Budget beer	2-4L	2.05	3.72	0.42
	4-6L	4.31	7.32	0.41
	8-20L	8.34	12.50	0.40
<i>Wine</i>				
Red wine	1x750ml	0.72	4.59	0.51
	2x750ml	1.22	7.90	0.50
	3x750ml	1.78	11.44	0.48
	4x750ml	3.08	18.35	0.47
White wine	1x750ml	0.72	4.44	0.51
	2x750ml	1.23	7.73	0.50
	3x750ml	1.76	10.99	0.49
	4x750ml	2.88	16.81	0.47
Rose wine	1x750ml	0.72	4.24	0.51
	2x750ml	1.79	10.14	0.48
Sparkling wine	1x750ml	0.74	5.24	0.71
	2x750ml	2.27	9.56	0.49
Champagne	1-2x750ml	1.16	26.54	1.86
Port	1-2x750ml	0.90	8.89	0.52
Sherry	1-2x750ml	1.20	7.85	0.41
Vermouth	1-2x750ml	1.32	7.12	0.39
Other fortified wines	1-2x750ml	1.33	6.63	0.36
<i>Spirits</i>				
Premium gin	1x700ml	0.69	11.74	0.45
	2x700ml	1.16	17.95	0.42
Budget gin	1x700ml	0.75	10.12	0.35
	2x700ml	1.27	14.89	0.33
Premium vodka	1x700ml	0.67	10.25	0.41
	2x700ml	1.16	16.17	0.39
Budget vodka	1x700ml	0.59	8.18	0.37
	2x700ml	1.14	14.71	0.36
Premium whiskey	1x700ml	0.67	19.60	0.72
	2x700ml	1.29	30.55	0.67
Budget whiskey	1x700ml	0.66	10.96	0.42
	2x700ml	1.21	16.17	0.40
Liqueurs; ABV <30%	1x700ml	0.64	7.82	0.68
	2x700ml	1.25	15.45	0.64
Liqueurs; ABV ≥30%	1x700ml	0.62	13.57	0.59
	2x700ml	1.14	21.38	0.52
Brandy	1x700ml	0.63	10.77	0.46
	2x700ml	1.11	17.43	0.44
Rum	1x700ml	0.79	12.26	0.41
	2x700ml	1.45	19.98	0.41
Pre-mixed spirits	700ml	0.70	4.37	0.94
Alcopops	1.3L	1.32	5.89	0.84
<i>Cider</i>				
Apple cider, <5% ABV	1L	0.91	2.66	0.65
	2-3L	2.45	3.94	0.37
	6-10L	7.10	10.15	0.37
Apple cider, 5-6% ABV	1-2L	1.62	2.77	0.32
	4L	3.79	5.49	0.27
	10-14L	8.41	10.39	0.28
Apple cider, >6% ABV	1-2L	1.21	3.28	0.38
	3-9L	4.40	7.06	0.27
Pear cider	1L	0.92	2.53	0.56
	3-6L	3.87	7.21	0.39
Fruit cider	750ml	0.68	2.48	0.81
	1-3L	1.98	6.75	0.75

Notes: Mean quantity is the average quantity of each product purchased by households in a given week over the calendar year. Mean price is the average price over regions and months in 2011.

Preference specification

We model consumer choice using a discrete choice demand model. On each week in which we observe a household purchasing groceries, we model which, if any alcohol variety the household chooses. The model rationalizes the fact that at the household-week level there are many zero demands for alcohol varieties. When we observe households purchasing more than one variety a week,¹⁹ we assume that these purchases are independent. This would be the case, if, for example, purchases are for different members of the household.

Let i index households, w index weeks and m index the “month” (based on 4 week periods) to which week w belongs. We model the decision household i in week w makes over what, if any, alcohol variety to purchase. We index the alcohol varieties $j = \{1, \dots, J\}$ and the decision to purchase no alcohol by $j = 0$. We assume the utility that household i in week w obtains from selecting alcohol variety $j > 0$ is given by:

$$U_{ijw} = -\alpha_i \rho_{jrm} + \mathbf{x}'_{jrm} \boldsymbol{\beta}_i + \epsilon_{ijw}, \quad (4.1)$$

where ρ_{jrm} is the price of variety j in region r and month m , \mathbf{x}_{jrm} are variety attributes (including number of alcoholic units, z_j), ϵ_{ijw} is a shock to utility that we assume is i.i.d. and $\boldsymbol{\theta}_i = (\alpha_i, \boldsymbol{\beta}_i)$ are preference parameters that govern how much weight the consumer’s utility function places on price and the various product attributes. We normalize utility from choosing no alcohol to $U_{i0w} = \epsilon_{i0w}$.²⁰

The vector of variety attributes includes alcoholic strength and a quadratic in total ethanol content, where we allow the first order term to vary with alcohol segment (beer, wine, spirits and cider). Preferences for these attributes capture the weight the household places on the strength and size of a variety, allowing the effect to differ depending on whether the variety is a beer, wine, spirit or cider. The product attributes also include a set of time varying alcohol type effects. These capture unobserved variation in preferences for alcohol types (e.g., conditional on alcohol content, a household may prefer gin to vodka), and the possibility these unobserved preferences fluctuate over time due, for instance, to seasonal patterns in demand.

We assume the i.i.d. shock to utility, ϵ_{ijw} , is distributed type I extreme value. This means the probability household i chooses variety $j > 0$ during week w , con-

¹⁹On 44% of the weeks in which households are observed purchasing alcohol they choose more than one variety; on more than 90% of weeks households are observed buying 3 or fewer varieties.

²⁰Suppressing household and time indices, the underlying utility problem is $V(\boldsymbol{\rho}, y, \mathbf{x}, \boldsymbol{\epsilon}) = \max_{(z, j \in \{0, \dots, J\})} \alpha z + \mathbf{x}'_j \boldsymbol{\beta}_i + \epsilon_{ij}$ subject to $z + \rho_j = y$, where z is consumption of a non-alcohol numeraire good and y is the consumer’s total budget. As y does not impact the choice of j it drops out of the discrete choice problem.

ditional on prices, product attributes and preferences takes the form:

$$\Pr(j|\rho_{jrm}, \mathbf{x}_{jm}, \boldsymbol{\theta}_i) = \frac{\exp(-\alpha_i \rho_{jrm} + \mathbf{x}'_{jm} \boldsymbol{\beta}_i)}{1 + \sum_{k>0} \exp(-\alpha_i \rho_{krm} + \mathbf{x}'_{km} \boldsymbol{\beta}_i)}. \quad (4.2)$$

We include in our demand model both posted price, ρ_j , and number of alcoholic units, z_j . As in Section 3 we define price per unit as $p_j = \rho_j/z_j$. We write the total number of units demanded by household i directly as a function of price per unit. In particular, household i 's total alcohol demand per adult per week during month m , when faced with prices per unit of alcohol, $\mathbf{p}_{rm} = (p_{1rm}, \dots, p_{Jrm})'$ is given by $Q_{im}(\mathbf{p}_{rm}) = \frac{1}{A_i} \frac{v_{im}}{4} \sum_j \Pr(j|\rho_{jrm}, \mathbf{x}_{jm}, \boldsymbol{\theta}_i) z_j$, where A_i is the number of adults in household i and v_{im} is the number of choices occasions on which we observe the household in month m .²¹

Preference heterogeneity

We incorporate two forms of preference heterogeneity into the model. We allow for all preference parameters to vary across nine household groups depending on whether the household contains light, moderate or heavy drinkers in the pre-sample period,²² and whether they have low, medium or high income.²³ In addition, we allow for unobserved household level preference heterogeneity.

The inclusion of preference heterogeneity by light, moderate and heavy drinkers is important for assessing the extent to which alcohol demands (and responses to policy reforms) vary across an important marker of the size of marginal externalities from drinking. We also model preference heterogeneity by income group to allow us to consider the equity implications of different policies – for instance, the extent to which a price floor or tax reform reduces the utility of low income households by more than higher income households.

²¹For a household that chooses either 0 or 1 alcohol variety each week, $v_{im} = \{1, 2, 3, 4\}$ and is given by the number of weeks in the month on which we observe them purchasing any groceries. However, sometimes we observe a household purchasing multiple varieties in a week (see footnote 19). In this case v_{im} can exceed 4. For instance, if in month m we observe a household purchasing 1 alcohol variety in each of the first, second and third weeks of the month, but 2 in the final, then $v_{im} = 5$.

²²We classify households into these three groups based on their average purchases of alcohol per adult per week in 2010 (the pre-sample period) – light is defined as less than 7 units, moderate as 7-14 units, and heavy as more than 14 units. These cutoffs are based on UK government recommendations.

²³We classify households into these groups based on the tercile of the distribution of annual household income per equivalized person.

We also allow for unobserved preference heterogeneity, which helps the model to capture realistic patterns of consumer switching across varieties.²⁴ We allow for household specific preferences for price, alcohol strength, ethanol content and the segment of the alcohol market to which the variety belongs. We model these effects as random coefficients that we assume are normally distributed within each of the nine groups. Total demand for alcohol among consumer group d , when faced with prices, \mathbf{p}_m is given by $Q_{dm} = \int Q_{im}(\mathbf{p}_m) dF_d(\boldsymbol{\theta}_i)$.

4.2 Identification

A central empirical challenge that we face is identifying the causal effect of price on demand. This requires us to isolate exogenous price variation (i.e., variation that is uncorrelated with demand factors for which we do not control). We do this by including a rich set of controls in the demand specification and using a control function and cost-based instruments in estimation.

Our demand controls include alcohol variety fixed effects, which capture unobserved quality differences across varieties. We allow for temporal variation in these effects, which captures fluctuations in national level demand, due, for instance, to seasonality and spikes in demand related to advertising campaigns. An important feature of the UK grocery industry is that the supermarkets have national price strategies (Competition Commission (2000)), which limits the scope for regional price setting driven by local demand shocks.

However, we cannot rule out the possibility that there may be some residual omitted demand side variables that are correlated with prices. We therefore include a control function that isolates price variation driven by a set of instruments that are likely to shift costs, but not to impact demand (Blundell and Powell (2004), Petrin and Train (2010)). These instruments include alcohol duty rates, producer price indices, exchange rates and oil prices. For details on the instruments and first stage, see Appendix B.1.

We allow for preference parameters to vary with a measure of households' long run alcohol consumption. However, we do not model high-frequency state dependence arising from the effect of recent past purchases conditional on current behavior. Current choices may depend on past choices due to high frequency habit formation, or due to households' stockpiling during sales periods. In previous work

²⁴See inter alia Berry et al. (2004) and Nevo (2001). In the absence of unobserved preference heterogeneity the model would exhibit the independence of irrelevant alternative property within drinking-household income groups cells. This would mean, for instance, for all households classified as light drinkers and low income, the cross-price elasticities associated with the change in price of a given variety would all be equal.

(Griffith, O’Connell, and Smith (2019)) we present evidence that these forms of dynamics are not of first order importance in the UK alcohol market, once we account for household level preference heterogeneity. Our longitudinal micro data helps us to identify a rich distribution of household preferences (see Berry and Haile (2020)).

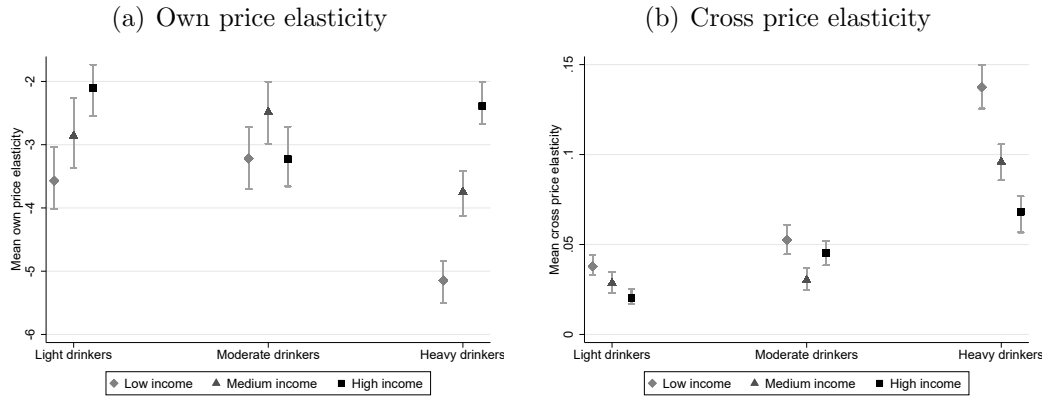
The combination of micro level data, a model with rich preference heterogeneity, and institutional features of the UK market gives us confidence in the validity of our demand estimates. Ultimately, though, we cannot directly test our identifying assumption that the residual price variation is uncorrelated with demand shocks. Instead, in Section 4.4 we provide evidence on the validity of our demand model by comparing the out-of-sample predictions of the model with the quasi-experimental estimates of behavior changes that we present in Section 3.4.

4.3 Elasticities

We estimate the model using simulated maximum likelihood. In Appendix B.2 we present the coefficient estimates from our model. The means of the distributions of preferences over price are all negative and statistically significant. There is statistically significant heterogeneity in preferences over price, alcohol strength, and ethanol quantity across households, even conditional on the nine household groups. We estimate the correlations between the preference parameters, which indicate that the more price sensitive households typically have stronger preferences over the quantity of ethanol and alcoholic strength of the variety.

The parameter estimates generate a set of own and cross price elasticities that describe how households switch between all the varieties in the market, as well as towards not buying alcohol, in response to marginal price changes. We summarize these in Figure 4.1. Panel (a) shows that the lowest income households, particularly the heavy drinkers, have the largest (in absolute terms) own-price elasticities – this indicates that they are the most willing to switch away from a particular variety if its price increases. Panel (b) shows the variation in cross-price elasticities: heavier drinkers and lower income households have larger cross-price elasticities – although they are more willing to switch away from a particular variety if its price changes, they are more likely to switch to another alcohol variety than to not buying alcohol.

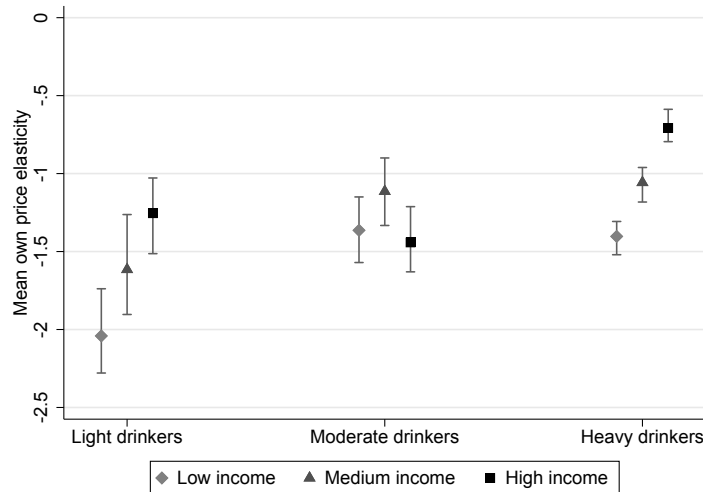
Figure 4.1: *Elasticities across types of drinkers and income levels*



Notes: Panel (a) shows the mean own-price elasticity across alcohol varieties for households in different groups and panel (b) shows the mean cross-price elasticity across variety pairs for households in different groups. In both figures, the bars illustrate 95% confidence intervals.

Figure 4.2 shows the percentage change in alcohol units demanded in response to a 1% change in the price of all alcohol varieties; this captures both households' substitution away from alcohol as well as any switching between different alcohol varieties that impacts their total ethanol demand. The figure shows that the most elastic households are low income, light drinkers, who have an own-price elasticity of all alcohol above 2 in absolute terms. At the other extreme are high income, heavy drinkers, whose elasticity is more than half as small in absolute terms.

Figure 4.2: *Total alcohol elasticity*



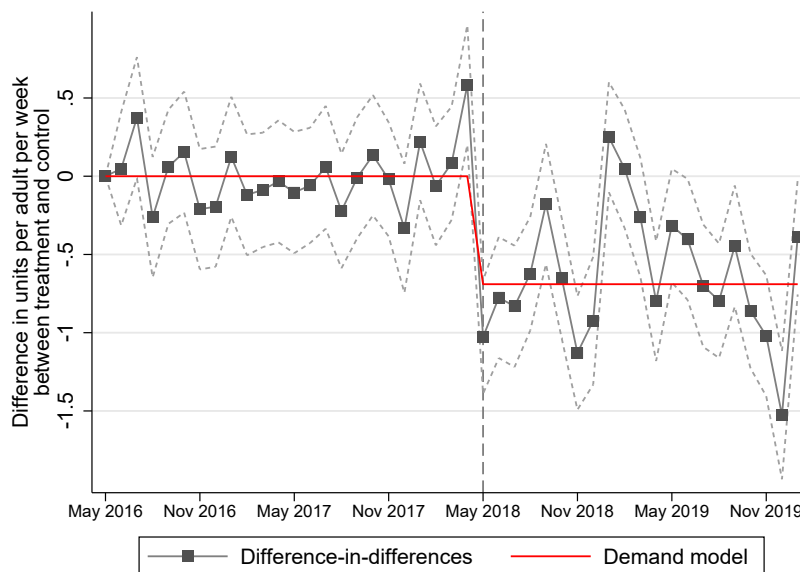
Notes: We simulate a 1% price increase of all alcohol varieties; the markers show the resulting percentage change in total units of alcohol demanded for the different household groups.

4.4 Model validation

A common concern with the approach we take to counterfactual policy analysis is that it imposes parametric restrictions that limit the credibility of our predictions. Here we show that our estimated demand model does a good job of predicting the effects of the price floor, obtained from our estimates based on the difference-in-differences approach in Section 3. Recall that our demand model is estimated on data from 2011, while the difference-in-difference analysis is based on a reform in 2018.²⁵

Mean effects Figure 4.3 compares the predicted average effect of the minimum unit price on alcohol purchases from our demand model with the quasi-experimental estimates (Figure 3.3). The model predicts that the minimum unit price leads to an average fall of 0.69 units per adult per week, which is not statistically different from the decline of 0.63 we estimate using the quasi-experimental variation.

Figure 4.3: *Predicted and observed impact of the minimum unit price, aggregate*



Notes: The figure repeats Figure 3.3 with the predictions from the demand model overlaid in red.

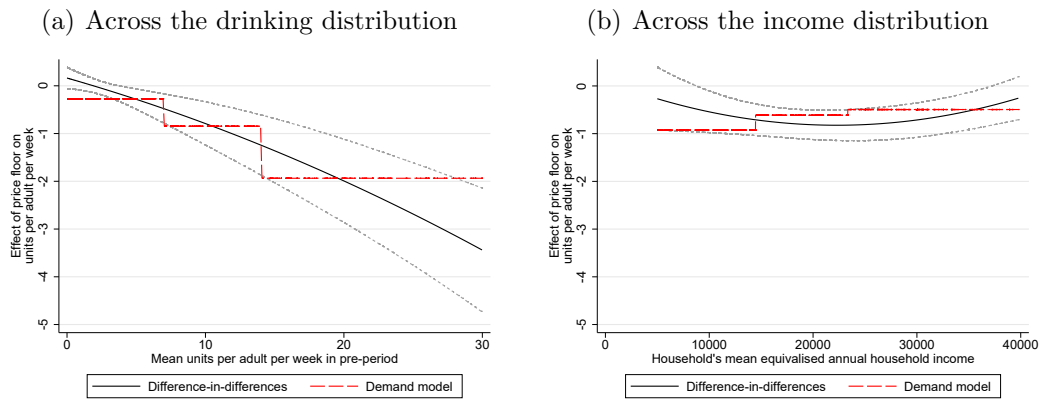
Heterogeneity across households. In Figure 4.4 we compare the demand model's predictions of the heterogeneous effect of the minimum unit price with the effects

²⁵To account for price inflation between 2011 (the time period we use to estimate demand) and 2018 (when the Scottish price floor was introduced), we simulate the introduction of a price floor that implies the same fraction of 2011 transactions are below the floor as the share in 2018 that are below a 50p/unit floor.

estimated using the quasi-experimental variation. Panel (a) shows how the impact of the minimum unit price on alcohol purchases varies by the average pre-period drinking of the household (the black line repeats the information in Figure 3.4). We categorize households into discrete groups to estimate the demand model, so the predictions are a step function. The figure shows that they match the quasi-experimental predictions well.

Panel (b) performs an analogous exercise showing variation across household income. We estimate a variant of equation (3.4), but replacing \bar{Q}_i with the household's mean equivalised annual income over the period. There is much less variation in the treatment effect with income than there is across the drinking distribution. The model's predictions of the effect lie with the 95% confidence bands of the impact of the minimum unit price estimated using the quasi-experimental variation.

Figure 4.4: *Predicted and observed impact of the minimum unit price: heterogeneity across households*

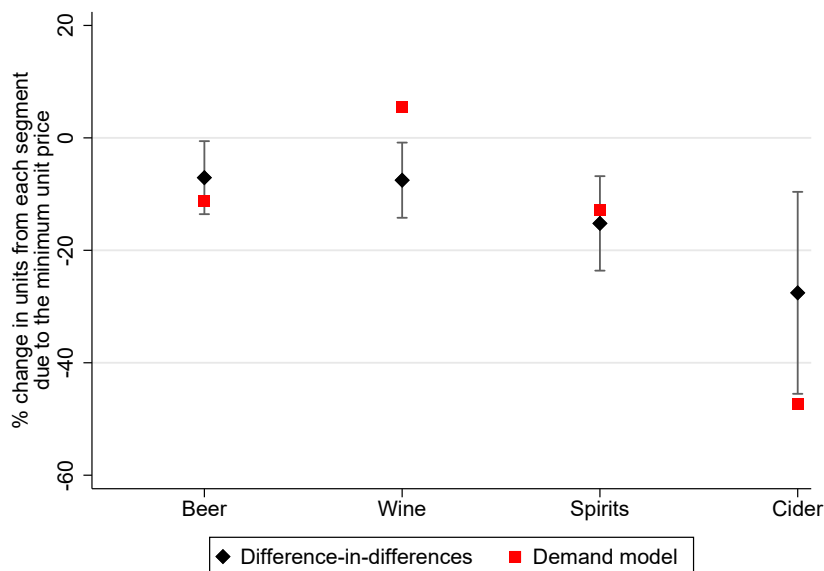


Notes: Panel (a) repeats Figure 3.4 with the predictions from the demand model overlaid in red. Panel (b) shows the analogous figure but where the treatment effect is allowed to vary by household income i.e., replacing \bar{Q}_i in equation (3.4) with the household's average equivalized annual income over the period.

Heterogeneity by alcohol type. We also compare the demand model's predictions and the quasi-experimental estimates of effects across alcohol types. Figure 4.5 shows the percent change in purchases for the alcohol segments (beer, wine, spirits and cider). The black markers show the estimated change in units from different alcohol types using the quasi-experimental variation and difference-in-differences approach described in Section 3. The largest percentage reductions are for cider, and the smallest for wine – this is intuitive given that cider saw the largest price increases and wine the lowest. The red markers show that the model predicts these

broad patterns, though with a larger predicted reduction for cider purchased, and a small increase (rather than decrease) in purchases of wine.

Figure 4.5: *Predicted and observed impact of the minimum unit price: heterogeneity across alcohol type*



Notes: The black markers show the percentage change in units from beer, wine, spirits, and cider estimated based on quasi-experimental variation and equation 3.3; the red markers show the predictions from the demand model.

5 The welfare effects of policy reforms

The demand model estimates allow us to compute the change in welfare under the price floor, and to compare this with alternative counterfactual policy reforms. We compare the impact of the price floor with a reform that replaces the existing system of alcohol duties with a single tax rate levied on the source of externalities (ethanol content).

5.1 Firm pricing response

We use the social welfare function described in Section 2. One of the assumptions we make is that the legal incidence of policy is fully reflected in prices. In practice, this means that the impact of a price floor is to increase the price of varieties that are priced below the floor to the floor and other prices remain unaffected, and that taxes are fully shifted to consumer prices. Whether this assumption is reasonable is an empirical question. Figure 3.1(c) shows that following the introduction of the

price floor in Scotland, the prices of products formerly below the floor moved to the floor, and there was little change in the prices of products above the floor.

We also provide evidence of complete pass-through of alcohol taxes to consumer prices in Appendix C.2. We use variation in the alcohol duty rates over the period 2010-12 and an approach similar to that of Besley and Rosen (1999). We find that there was no statistically significant change in the *tax-exclusive* price as a result of duty changes, and that the consumer prices changed by one plus the VAT rate. These results are consistent with complete pass-through of alcohol taxes to prices.

5.2 External costs and welfare changes

For reference, the social welfare function set out in Section 2 that we use to evaluate the welfare impact of different policies is:

$$W(\mathcal{P}^\iota) = \sum_i v_i(\mathbf{p}(\mathcal{P}^\iota)) + \sum_j \tau Q_j(\mathcal{P}^\iota) - \sum_i \Psi_i(Q_i(\mathcal{P}^\iota)), \quad (5.1)$$

where $v_i(\cdot)$ is the money metric utility of consumer i ,²⁶ \mathbf{p} is the vector of alcohol prices (expressed per unit), Q_j is total demand in units for variety j , and $\Psi_i(\cdot)$ the total externality from alcohol consumption of consumer i (we discuss the form this takes below). Here we write prices and quantities as a function of \mathcal{P}^ι , which denotes the policy environment.

Policy environments

We denote the baseline policy environment by \mathcal{P}^0 . This consists of the alcohol duties that were in place in the UK in 2011, which vary in both rates and bases across alcohol segment (beer, wine, spirits and cider) – see Appendix C.1 for a description of the tax system. In 2011 there was no price floor.

We compare the effect of two policy reforms: \mathcal{P}^1 denotes the baseline policy environment plus a price floor, and \mathcal{P}^2 denotes a reform that replaces the baseline system of alcohol duties with a single rate levied in proportion to ethanol content. We measure their impact relative to the baseline policy environment, i.e., we compare $W(\mathcal{P}^1) - W(\mathcal{P}^0)$ and $W(\mathcal{P}^2) - W(\mathcal{P}^0)$. Below we show robustness of our comparisons of the price floor and tax reform to this baseline environment.

²⁶Given our demand model this is given by the expression $v(\rho_{jr(i)m(w)}, \mathbf{x}_{jm(w)}, \boldsymbol{\theta}_i) = \ln \sum_{j>0} \exp(-\alpha_i \rho_{jr(i)m(w)} + \mathbf{x}'_{jm(w)} \boldsymbol{\beta}_i)$.

Denote by p_{jrm}^ι the price per unit of variety j in region r and month m under policy environment ι . Prices under the three policies can be written:

$$\begin{aligned} p_{jrm}^0 &= (w_{jrm} + \tilde{\tau}_j^0)(1 + \tau_{VAT}) \\ p_{jrm}^1 &= \max\{\underline{p}, (w_{jrm} + \tilde{\tau}_j^0)(1 + \tau_{VAT})\} \\ p_{jrm}^2 &= (w_{jrm} + \tau)(1 + \tau_{VAT}), \end{aligned}$$

where w_{jrm} denotes the variety's tax exclusive price per unit, $\tilde{\tau}_j^0$ is the amount of alcohol duty per unit levied on the variety under the observed tax system, \underline{p} denotes the price floor (under policy environment 1) and τ is the ethanol tax rate (under policy environment 2). We assume that the system of value added taxation is unchanged across the three policy environments; we show robustness of our results to a zero-VAT environment below.

We choose the price floor \underline{p} to be equivalent to level set in Scotland in 2018, and, to make the ethanol tax comparable, we choose the tax rate that leads to the same reduction in aggregate alcohol units purchased as under the price floor.²⁷

Externality function

We assume that the externality function, $\Psi_i(\cdot)$, is an increasing weakly-convex function of the form:

$$\Psi(Q_i) = A Q_i^b, \text{ where } A > 0 \quad b \geq 1. \quad (5.2)$$

When $b = 1$ the function is linear in ethanol and the marginal externality is given by A and is common across each unit of alcohol consumed. When $b > 1$ the function is convex and the marginal externality is increasing in alcohol consumed.

We show how the impact of policy reforms varies with the degree of convexity of the externality function, holding fixed the total external costs created by drinking. The parameter b determines the degree of convexity of the function ($(\Psi_i'' Q_i)/\Psi_i' = b - 1$): we vary b between 1 and 4. To make this interpretable, we express this in terms of the ratio of the marginal externality of the drinker at the 90th percentile of the drinking distribution to that of the median drinker. For instance, when $b = 1$, the externality function is linear, which means that the marginal externalities of the 90th percentile and median drinker are the same. At the other extreme, when $b = 4$, the marginal externality of the 90th percentile drinker is 125 times as large

²⁷For notational parsimony, we suppress the month, m , subscripts on the social welfare function (equation (5.1)). In practice, we simulate the effect of the policy reforms in each month using the average price across regions and then aggregate up to the yearly level to get a measure of annual changes in welfare.

as the marginal externality of the median drinker. Table C.2 lists this ratio for all calibrations. Conditional on b , we calibrate A so that the total external costs under the baseline policy environment are equal the central estimate used by the UK government;²⁸ below we show the robustness of our results to varying the total external costs from drinking.

Components of welfare

Table 5.1 summarizes the impact of the price floor and ethanol tax on welfare under a linear externality function (when $b = 1$ in equation (5.2)). The first three rows of the table show the change in the components of the social welfare function. The ethanol tax lowers consumer surplus by more than twice as much as the price floor. However, it raises tax revenue by £565 million, whereas the price floor leads to a £371 million decline in tax revenue. By construction, under the linear externality calibration, both policies achieve the same reduction (of £639 million) in the external costs of drinking. The fourth row shows the impact of each policy on social welfare. The ethanol tax leads to a £198 million rise in welfare, whereas a price floor lowers welfare by £225 million. The fifth row shows the impact of the two policies on firms' profits (under the assumption that the market is competitive). The price floor confers windfall profits of £293 million. If these are included in social welfare, the price floor results in an increase of £68m, which is still lower than the £198 million increase under the ethanol tax.

When externalities are linear an ethanol tax that leads to the same reduction in external costs will always outperform a price floor. However, as discussed in Section 3.1, there is evidence that externalities from alcohol are convexly increasing in total intake and thus vary across households. In this case an ethanol tax is no longer necessarily preferred and it will depend on the pattern of demand responses across the drinking distribution; a price floor may perform better.

²⁸We use the estimate of the direct tangible social costs in Cnossen (2007), which is based on a study by the UK Cabinet Office (2003). Accounting for price inflation and the fact that we are studying alcohol purchases not made in restaurants and bars gives a total external cost of £7.25 billion per year.

Table 5.1: *Effect of the price floor and ethanol tax under linear externalities*

Change, relative to baseline, \mathcal{P}^0 , under:	Price floor, \mathcal{P}^1 (£m/year)	Ethanol tax, \mathcal{P}^2 (£m/year)
Consumer surplus	-494	-1006
Tax revenue	-371	565
External costs (under linear ext. function)	-639	-639
Social welfare (under linear ext. function)	-225	198
Firm profits	293	0
Social welfare + profits (under linear ext. function)	68	198

Notes: Social welfare is equal to the sum of consumer surplus, tax revenue minus the external costs from alcohol consumption. External costs are calculated using the externality function (5.2), with $b = 1$. Tax revenue includes revenue from all duties applied to alcohol and revenue from VAT on alcohol purchases. Numbers are grossed up to the annual level for the UK as a whole. Profits are calculated assuming that the market is perfectly competitive and tax-exclusive prices are equal to marginal cost.

Table 5.2 compares the impact of a price floor and the ethanol tax on alcohol units purchased per adult per week across light, moderate and heavy drinkers. It shows that the price floor leads to a 40% larger reduction in the units purchased by the heavy drinkers than the ethanol tax. Conversely, the reduction in units purchased by the light drinkers is a third smaller under the price floor than under the ethanol tax.

Table 5.2: *Heterogeneous impact of the price floor and ethanol tax*

	Change in units per adult per week		% of aggregate decline from group	
	Floor	Tax	Floor	Tax
Light drinker	-0.27	-0.42	25.4	39.3
Moderate drinker	-0.85	-0.91	20.9	22.4
Heavy drinker	-1.94	-1.38	53.7	38.3
Total	-0.69	-0.69	100.0	100.0

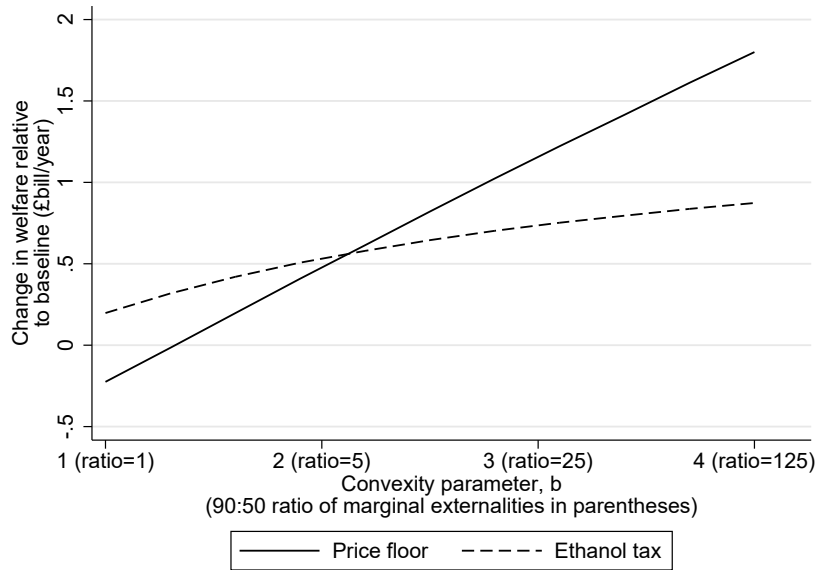
Notes: Change in units per adult per week is relative to the UK tax system with no price floor in place, \mathcal{P}^0 .

When the externality function is convex, these differences in how policy induces change in purchases across the drinking distribution change the relative performance of the two policies. To show this, we calculate the change in social welfare under the alternative policy reforms when we adjust the convexity of the externality function (5.2). As described above, we vary b between 1 and 4, and adjust A so that the total external costs under the baseline policy environment are constant across

calibrations: this allows us to isolate the impact of the degree of convexity, rather than magnitude, of externalities.

Figure 5.1 shows the difference in social welfare between the ethanol tax and price floor as we vary b . On the horizontal axis in parentheses, for each value of b , we report the ratio of the marginal externality of a drinker at the 90th percentile to the 50th percentile of the drinking distribution. When the marginal externality of a drinker at the 90th percentile is less than 7 times the marginal externality of the median drinker, the ethanol tax performs better. When the ratio is higher, the price floor leads to larger welfare gains than the ethanol tax: its superior targeting properties outweigh the fact that it leads to a reduction in tax revenue and a transfer of surplus to producers. When the 90:50 ratio of marginal externalities equals 30 – which is equivalent to heavy drinkers creating 85% of the total external costs under the baseline policy environment – the price floor leads to £470 million higher welfare than the ethanol tax.

Figure 5.1: *Difference in welfare under a price floor and ethanol tax*



Notes: We calculate the change in social welfare under the two different policy environments – the price floor and ethanol tax – relative to the baseline, under different calibrations of the externality function (equation (5.2)). We vary b between 1 and 4, and adjust A so that the total externality under the baseline policy environment is constant across calibrations. The horizontal axis the convexity, measured as $b = (\Psi_i'' Q_i) / \Psi_i' + 1$. Welfare is measured in £billion per year for the UK.

5.3 Robustness

Baseline tax system

The impact of a price floor depends, to a certain extent, on the tax system under which it is implemented. In the analysis above, we consider a price floor introduced on top of the existing tax system, as has been done in practice in Scotland. Here we consider the implications of varying the baseline tax system for our results. In each case we use the same externality calibration as above.²⁹

First, we consider a counterfactual world in which the UK tax rates are all reduced by 50%. This is motivated by the fact that, although the variation in rates across alcohol segments is similar across countries, average rates vary considerably. For example, alcohol taxes in the US are substantially lower than in the UK. We consider a baseline policy environment with 50% lower tax rates, and we study the impact of a price floor and an ethanol tax that lead to an aggregate reduction in units purchased of 0.69 units per adult per week relative to this baseline. Figure 5.2(a) shows how the impact of the price floor and ethanol tax (relative to this baseline environment) vary with the degree of convexity of the externality function. Consistent with the results outlined in the previous section, under linear externalities the ethanol tax outperforms the price floor, but if the externality function is sufficiently convex the price floor leads to larger welfare gain than the ethanol tax (with the gap between the two policies increasing in the degree of externality convexity).

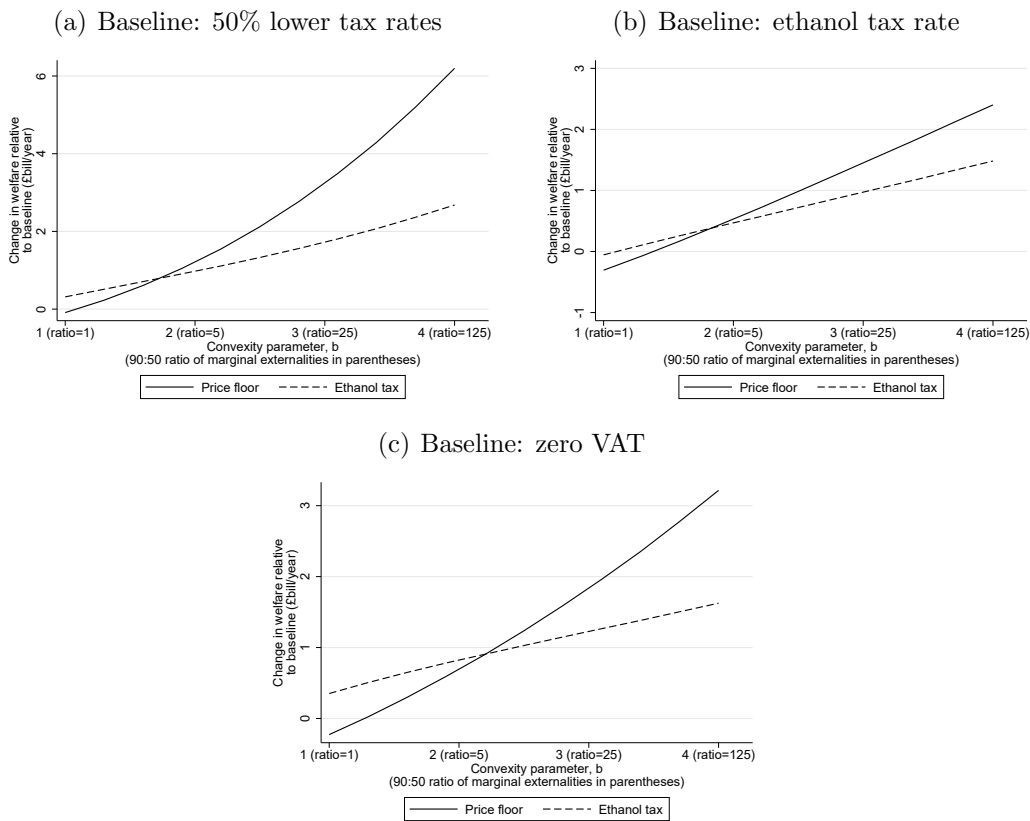
Second, we consider a scenario in which the baseline policy environment is one with an ethanol tax and we investigate whether increasing the ethanol tax rate or introducing a price floor leads to larger welfare gains. As the baseline, we choose an ethanol tax rate such that the aggregate units purchased are the same as under the existing UK system. We then consider a price floor and ethanol tax rate increase that leads to an aggregate reduction in units purchased of 0.69 units per adult per week relative to this baseline. Figure 5.2(b) shows that when the externality function is sufficiently convex, introducing a price floor is preferable to increasing the ethanol tax rate.

Third, we undertake the same exercise as in the previous section after first removing the value-added tax (VAT). In this case the baseline policy environment

²⁹Specifically, when varying the convexity of the externality function (i.e. b), we use the same corresponding values of A as in the preceding section. Therefore, here we focus on robustness of our results to the baseline policy environment, holding the aggregate externality for any given value of b and distribution of drinking fixed. In the next section we assess robustness to the size of aggregate externalities.

corresponds to the UK system of alcohol duties, but with no VAT. Again we simulate the effect of introducing a price floor and of replacing the alcohol duties with an ethanol tax, both set so they lead to an aggregate reduction in units purchased of 0.69 units per adult per week relative to the baseline. Figure 5.2(c) shows that, as above, for a sufficiently convex externality function the price floor outperforms the ethanol tax.

Figure 5.2: *Difference in welfare under minimum unit price and ethanol tax under alternative baseline tax systems*



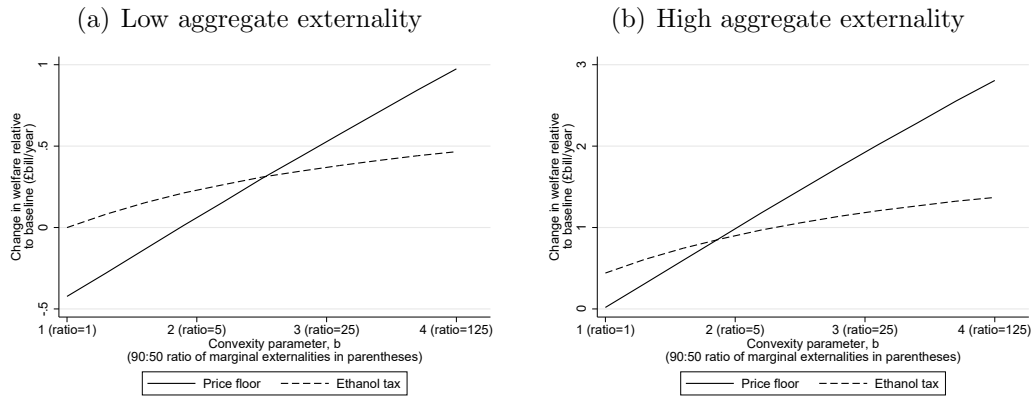
Notes: Each panel shows how the impact of introducing a price floor and of replacing the baseline tax system with an ethanol tax varies with the degree of convexity of the externality function. In all cases, the price floor and ethanol tax rate are set achieve a 0.69 reduction in units purchased per adult per week relative to the new baseline policy environment. The baseline policy environment varies across panels and is described in the text. We vary b between 1 and 4, and adjust A so that the total external costs under the UK tax system (with no price floor) are constant across calibrations. Welfare is measured in £ billion per year for the UK.

Total external costs

In our central calibration, we set A and b such that the total external costs of alcohol consumption are equal to the UK government's estimate of the external costs of drinking (£7.25 billion). In Figure 5.3 we show how our results vary when

we re-calibrate A so that the total external costs are lower at £5 billion (left hand panel), or higher at £10 billion (right hand panel). As in our central calibration, the price floor outperforms the ethanol tax when the externality function is sufficiently convex. The higher are total externalities, the lower the degree of convexity of the externality function necessary for the price floor to outperform the ethanol tax.

Figure 5.3: *Difference in welfare under minimum unit price and ethanol tax under higher and lower total externalities*

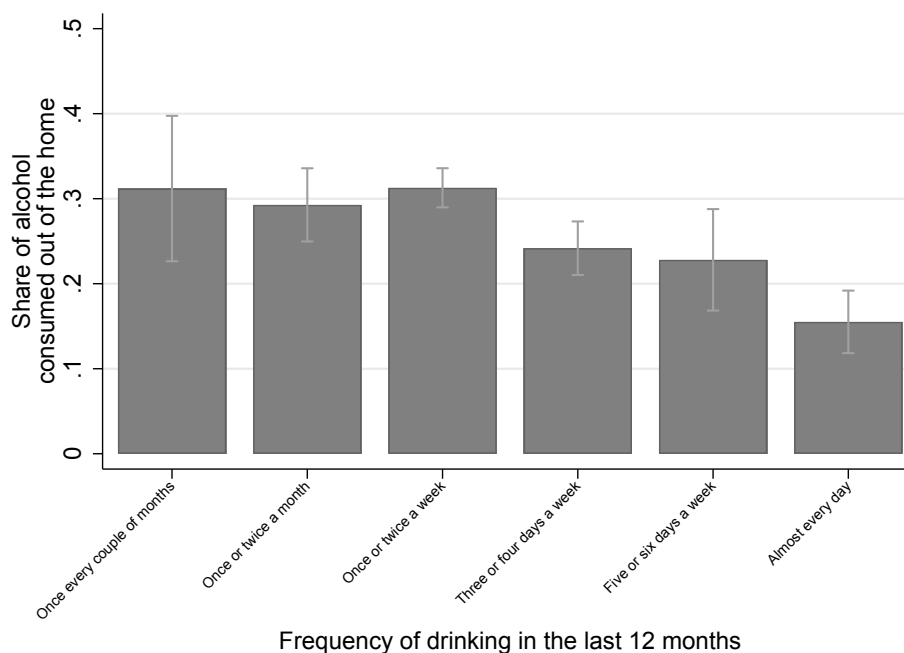


Notes: The left hand panel shows the social welfare under an alternative calibration of the externality function in which the total external costs under the UK tax system (with no price floor) are equal to £5 billion, and the right hand panel shows results when the total external costs are £10 billion. Welfare is measured in £ billion per year for the UK.

Alcohol consumed out of the home

Our data covers alcohol consumed at home, which makes up around three-quarters of units consumed in the UK. A price floor and ethanol tax are likely to affect the prices of alcohol consumed in restaurants and bars differently: in general, alcohol consumed out is priced above the floor and therefore unlikely to experience a price increase. As discussed in Section 3.1, the Scottish Government find no difference in the change from 2017 to 2018 in aggregate purchases of alcohol in restaurants and bars between Scotland and England. Figure 5.4 shows that heavier drinkers consume a greater fraction of their alcohol at home. This suggests that omitting alcohol consumed out is likely to underestimate the difference in welfare achieved under a price floor compared with an ethanol tax, because the former targets a greater fraction of total alcohol consumption (from all sources) for heavier drinkers.

Figure 5.4: *Alcohol consumed out of the home, by drinking frequency*



Notes: Figure drawn using data from the National Diet and Nutrition Survey. Data on share of alcohol consumption outside the home recorded using an intake diary; data on frequency of drinking over the past twelve months collected using a survey.

5.4 Distributional implications

So far we have abstracted from distributional considerations and focused on the efficiency of the two policies. Here we show the distributional effects of the policy reforms. This could, for example, inform any adjustments to a nonlinear income tax system that would be necessary to offset the distributional consequences of reform.

Table 5.3 summarizes the change in consumer surplus experienced by the different income groups under the price floor and ethanol tax. A common criticism of the price floor in policy debates is that it is regressive and places an unfair burden on lower income households. The table indicates that it is indeed the case that a greater *proportion* of the overall reduction in consumer surplus is borne by lower income households under the price floor (40%) when compared with the ethanol tax (33%). However, the aggregate reduction in consumer surplus is less than half as large as under the ethanol tax. This is because the price increases under the price floor are better targeted at reducing ethanol purchases, and so a smaller average price rise is required to achieve the same aggregate reduction in units per adult per week. The fall in consumer surplus for the low income households is therefore still less than under the ethanol tax.

Table 5.3: *Change in consumer surplus across the income distribution*

	(1)		(2)		(3)		(4)	
	Price floor				Ethanol tax			
	£m/year	% total	£m/year	% total	£m/year	% total	£m/year	% total
Low income	-196	40	-337	34	-337	34	-337	34
Medium income	-157	32	-315	31	-315	31	-315	31
High income	-141	28	-354	35	-354	35	-354	35
All households	-494	100	-1006	100	-1006	100	-1006	100

Notes: Columns (1) and (3) show the reduction in consumer surplus for all households within each income group (£million per year). Columns (2) and (4) express this as a percentage of the overall decline in consumer surplus.

As long as the price floor or ethanol tax does not lower overall social welfare, in principle households can be fully compensated for consumer surplus losses due to the policy from the savings from externalities net of changes in alcohol tax revenue. Figure 5.1 shows that this is true for all degrees of convexity of the externality function for the ethanol tax, and feasible for the price floor as long as the marginal externality created by someone at the 90th percentile is more than 60% larger than that of the median drinker. Table 5.3 indicates that the group of low income households would have to receive a proportionately greater share of the compensation under the price floor than under the ethanol tax. In practice, any such compensation would likely be carried out by adjustment to income taxation, which could be used to compensate consumers according to their incomes, but would be unable to compensate consumers with different alcohol consumption patterns, conditional on income.

Our analysis assumes that the planner does not aim to use alcohol taxation as a means to redistribute across the income distribution. Saez (2002) and Allcott et al. (2019) show that if there is preference heterogeneity for a good that is correlated with incomes then, even with recourse to a non-linear income tax, there is a motive to redistribute through the commodity taxation. We leave these considerations to future work.

6 Summary and conclusions

In this paper, we compare the relative effectiveness of a price floor with an ethanol tax for reducing the externalities generated by alcohol consumption. We exploit a natural experiment that involved the introduction of a price floor in one part of the UK (Scotland), but not in other parts, and show that the policy achieved greater

reductions in units of alcohol purchased amongst the heaviest drinkers. It is well targeted because heavy drinkers get a disproportionate amount of their alcohol from cheap products and because their demand for these products is relatively elastic.

We estimate a model of demand for alcohol products based on data prior to the announcement of a price floor and show the model's out-of-sample predictions match our estimates of the effects of the price floor based on the quasi-experimental variation. We use the model to assess the effects of the price floor on welfare, and to simulate a counterfactual ethanol tax, which is levied directly on the source of externalities. We show that if the externality from another drink is equal across all drinkers the ethanol tax outperforms the price floor. However, if externalities are at least moderately convexly increasing in intake – i.e., at the margin the consumption of heavy drinkers is more socially costly than light drinkers – the superior targeting properties of the price floor outweigh the fact that it leads to a transfer from consumers and government revenue to firm profits.

Our focus is on tax reform that replaces the existing system of alcohol duties with an ethanol tax. In our previous work (Griffith, O'Connell, and Smith (2019)) we show that when externalities are heterogeneous, then there are welfare gains from varying tax rates across products. The idea is related to that described here: correlation between the demands for different types of products and marginal externalities provide a way to tag socially costly consumption. A promising avenue for future research is to further consider how policy instruments such as taxation and price floors may be effectively combined to deal with externalities created by a range of goods.

References

- Allcott, H., B. B. Lockwood, and D. Taubinsky (2019). Regressive Sin Taxes, with an Application to the Optimal Soda Tax. *Quarterly Journal of Economics* 134(3), 1557–1626.
- Angrist, J. D. and J.-S. Pischke (2010). The Credibility Revolution in Empirical Economics: How Better Research Design is Taking the Con out of Econometrics. *Journal of Economic Perspectives* 24(2), 3–30.
- Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market. *Journal of Political Economy* 112(1), 68–105.
- Berry, S. T. and P. A. Haile (2020). Nonparametric Identification of Differentiated Products Demand Using Micro Data. *Cowles Foundation Discussion Paper No. 1718*, 29.
- Besley, T. J. and H. S. Rosen (1999). Sales Taxes and Prices: An Empirical Analysis. *National Tax Journal* 52(2), 157–78.
- Blundell, R. W. and J. L. Powell (2004). Endogeneity in semiparametric binary response models. *Review of Economic Studies* 71(3), 655–679.
- Centers for Disease Control and Prevention (2016). CDC - Data and Maps - Alcohol. <http://www.cdc.gov/alcohol/data-stats.htm>.
- Chetty, R. (2009). Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. *Annual Review of Economics* 1(1), 451–488.
- Cnossen, S. (2007). Alcohol taxation and regulation in the European Union. *International Tax and Public Finance* 14(6), 699–732.
- Competition Commission (2000). *Supermarkets: A Report on the Supply of Groceries from Multiple Stores in the United Kingdom*. London, U.K.: The Stationery Office.
- Conlon, C. T. and N. L. Rao (2019). The Price of Liquor is Too Damn High: Alcohol Taxation and Market Structure. *NYU Wagner Research Paper No. 2610118*, 58.
- Court of Justice of the European Union (2015). Scotch Whisky Association and Others v Lord Advocate. Technical report.
- Diamond, P. A. (1973). Consumption externalities and corrective imperfect pricing. *The Bell Journal of Economics* 4(2), 526–538.
- Dubois, P., R. Griffith, and M. O’Connell (2020). How Well Targeted Are Soda Taxes? *American Economic Review* forthcoming.
- Gillitzer, C., H. J. Kleven, and J. Slemrod (2015). A Characteristics Approach to Optimal Taxation: Line Drawing and Tax-Driven Product Innovation. *Scandinavian Journal of Economics*.
- Griffith, R., M. O’Connell, and K. Smith (2019). Tax design in the alcohol market. *Journal of Public Economics* 172, 20–35.
- Grigolon, L., M. Reynaert, and F. Verboven (2018). Consumer valuation of fuel costs and tax policy: Evidence from the European car market. *American Economic Journal. Economic Policy* 10(2), 193–225.
- Hamilton, A. (1790). First Report on the Public Credit. Technical report.

- Jacobsen, M. R., C. R. Knittel, J. M. Sallee, and A. A. van Benthem (2020). The Use of Regression Statistics to Analyze Imperfect Pricing Policies. *Journal of Political Economy* 128(5), 1826–1876.
- Lönnroth, K., B. G. Williams, S. Stadlin, E. Jaramillo, and C. Dye (2008). Alcohol use as a risk factor for tuberculosis—a systematic review. *BMC Public Health* 8(1).
- Miravete, E. J., K. Seim, and J. Thurk (2018). Market Power and the Laffer Curve. *Econometrica* 86(5), 1651–1687.
- Miravete, E. J., K. Seim, and J. Thurk (2020). One Markup to Rule Them All: Taxation by Liquor Pricing Regulation. *American Economic Journal: Microeconomics* 12(1), 1–41.
- National Health Service (2018). The risks of drinking too much. <https://www.nhs.uk/live-well/alcohol-support/the-risks-of-drinking-too-much/>.
- Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69(2), 307–342.
- NHS Health Scotland (2019). Monitoring and Evaluating Scotland’s Alcohol Strategy. Technical report.
- Petrin, A. and K. Train (2010). A Control Function Approach to Endogeneity in Consumer Choice Models. *Journal of Marketing Research* 47(1), 3–13.
- Rehm, J., B. Taylor, S. Mohapatra, H. Irving, D. Baliunas, J. Patra, and M. Roerecke (2010). Alcohol as a risk factor for liver cirrhosis: A systematic review and meta-analysis. *Drug and alcohol review* 29(4), 437–445.
- Saez, E. (2002). The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes. *Journal of Public Economics* 83(2), 217–230.
- Scottish Government (2018). Minimum Unit Pricing of Alcohol Consultation Document. Technical report.
- Seiler, S., A. Tuchman, and S. Yao (2020). The Impact of Soda Taxes: Pass-through, Tax Avoidance, and Nutritional Effects. *Journal of Marketing Research forthcoming*.
- Seim, B. K. and J. Waldfogel (2013). Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Board’s Entry Decisions. *American Economic Review* 103(2), 831–862.
- Stockwell, T., M. C. Auld, J. Zhao, and G. Martin (2012, May). Does minimum pricing reduce alcohol consumption? The experience of a Canadian province. *Addiction* 107(5), 912–920.
- Stockwell, T., J. Zhao, N. Giesbrecht, S. Macdonald, G. Thomas, and A. Wettlaufer (2012). The raising of minimum alcohol prices in Saskatchewan, Canada: Impacts on consumption and implications for public health. *American Journal of Public Health* 102(12), e103–e110.
- Stockwell, T., J. Zhao, A. Sher, R. C. Callaghan, S. Macdonald, and J. Gatley (2017, July). Assessing the impacts of Saskatchewan’s minimum alcohol pricing regulations on alcohol-related crime. *Drug and Alcohol Review* 36(4), 492–501.
- UK Cabinet Office (2003). Alcohol misuse: How much does it cost? Technical report.
- World Health Organization (2006). Intimate partner violence and alcohol. *Fact Sheet*.

- World Health Organization (2014). Global Status Report on Alcohol and Health. Technical report, Geneva: World Health Organization.
- World Health Organization (2016). *The Health and Social Effects of Nonmedical Cannabis Use*.
- World Health Organization (2017). Resource tool on alcohol taxation and pricing policies. Technical report.
- Zhao, J. and T. Stockwell (2017). The impacts of minimum alcohol pricing on alcohol attributable morbidity in regions of British Columbia, Canada with low, medium and high mean family income. *Addiction* 112(11), 1942–1951.

APPENDIX

FOR ONLINE PUBLICATION

Price floors and externality correction

Rachel Griffith, Martin O'Connell and Kate Smith

A Data

A.1 Comparison of Kantar with LCFS

Table A.1: *Household characteristics*

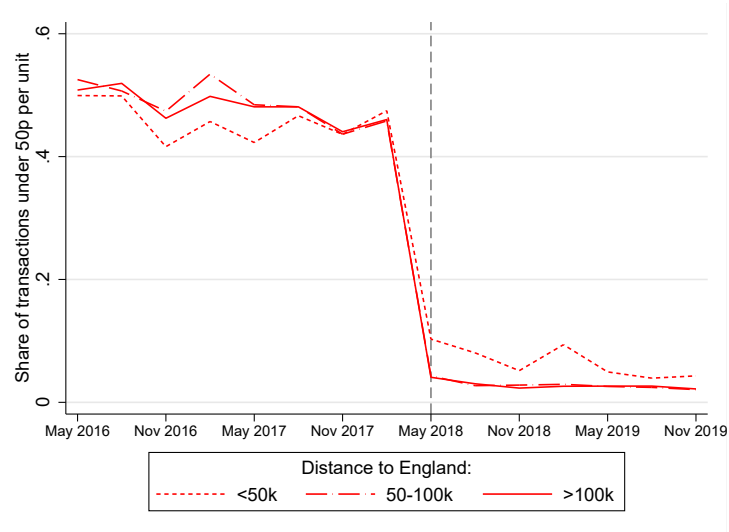
	Kantar (2011)	Kantar (2016)	LCFS
<i>Region</i>			
England - North	28.0 [27.2, 28.8]	28.0 [27.4, 28.5]	27.4 [26.1, 28.6]
England - Midlands	18.9 [18.2, 19.7]	18.2 [17.7, 18.7]	19.1 [18.0, 20.2]
England - South and East	44.5 [43.6, 45.4]	45.2 [44.6, 45.8]	44.4 [43.0, 45.8]
Scotland	8.6 [8.0, 9.1]	8.6 [8.3, 9.0]	9.1 [8.3, 9.9]
<i>Employment status of household head</i>			
Full time	38.6 [37.7, 39.5]	41.7 [41.1, 42.3]	39.6 [38.2, 41.0]
Part time	18.1 [17.4, 18.8]	20.9 [20.4, 21.4]	11.1 [10.2, 12.0]
Self-employed*	[., .]	[., .]	7.9 [7.2, 8.7]
Unemployed	2.3 [2.0, 2.6]	1.9 [1.8, 2.1]	2.4 [1.9, 2.8]
Retired or not working	40.9 [40.0, 41.9]	35.4 [34.8, 36.0]	39.0 [37.6, 40.4]
<i>Socioeconomic status</i>			
Highly skilled	22.1 [21.2, 23.0]	21.8 [21.2, 22.4]	18.7 [17.2, 20.1]
Semi skilled	60.9 [59.9, 62.0]	59.7 [59.0, 60.4]	59.4 [57.6, 61.2]
Unskilled	17.0 [16.1, 17.8]	18.4 [17.9, 19.0]	21.9 [20.4, 23.4]
<i>Alcohol purchases</i>			
Average weekly units, conditional on buying	14.9 [14.5, 15.4]	12.9 [12.6, 13.1]	14.5 [13.8, 15.1]

Notes: Table shows the share of households in the Kantar FMCG Purchase Panel and Living Costs and Food Survey (LCFS) in various demographic groups. Numbers are shown for the most recently available data for the LCFS, which is 2014, and for 2011 (the year we use to estimate demand) and May 2016 – April 2017 (the first year of our difference-in-differences analysis) in the Kantar data. *The self-employed are not distinguished from employees in the Kantar data. Socioeconomic status is based on the occupation of the head of the household and is shown for the set of non-pensioner households. Average weekly units are measured per adult, are for at-home consumption only, and are conditional on the household buying alcohol across the two week measurement period. 95% confidence intervals are shown below each share.

A.2 Treatment and control areas

A price floor for alcohol was introduced in Scotland in May 2018. A possible concern with comparing prices and purchases in Scotland and England is the conflating effect of cross border shopping. For example, Figure A.1 shows that there were about twice as many transactions below the floor in the region of England within 50km from the Scottish-English border than further away, following the introduction of the price floor. We therefore exclude a 50km region on either side of the Scottish-English border from our treatment and control groups.

Figure A.1: *Share of transactions below the price floor, by distance from the border*



Notes: The line shows the share of transactions under the £0.50 unit price before and after the introduction of the price floor in May 2018 for households that live < 50, 50 – 100 and > 100km from the Scottish-English border.

A.3 Difference-in-differences

Table A.2: Summary of aggregate results for units purchased per adult per week

	(1) Total	(2) Above floor	(3) Below floor	(4) Beer	(5) Wine	(6) Spirits
$treat_i \times post_w$	-0.631 (0.124)	0.258 (0.060)	-0.935 (0.116)	-0.215 (0.058)	-0.192 (0.087)	-0.225 (0.063)
Percentage change in units	-11.0	14.6	-27.6	-12.7	-7.5	-15.2
Time effects	Yes	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The first row shows the coefficient on the interaction of the $treat_i$ dummy (equal to 1 if the household is located in Scotland) and the $post_w$ dummy (equal to 1 if the transaction took place after the introduction of the price floor in Scotland). Column (1) shows the results for total alcohol purchases, (2) for products priced above the floor pre-reform, (3) for products previously priced below the floor, (4) for beer, (5) for wine, and (6) for spirits. All regressions include household fixed effects, year-month effects and controls for major holidays.

Table A.3: Summary of heterogeneity results for units purchased per adult per week

	(1) Total	(2) Above floor	(3) Below floor	(4) Total
$treat_i \times post_w$	0.163 (0.105)	0.045 (0.064)	0.431 (0.108)	0.099 (0.490)
$treat_i \times post_w \times Q_i$	-0.083 (0.037)	0.032 (0.021)	-0.176 (0.039)	
$treat_i \times post_w \times Q_i^2$	-0.001 (0.001)	-0.000 (0.000)	-0.001 (0.001)	
$treat_i \times post_w \times Y_i$				-0.000 (0.000)
$treat_i \times post_w \times Y_i^2$				0.000 (0.000)
Time effects	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes

Notes: The first three columns show the heterogeneity in treatment effect by the household's average units purchased in the pre-sample period, Q_i : for total alcohol (column 1), units priced above the floor pre-reform (2), and products previously priced below the floor (3). Column (4) shows the heterogeneity in the treatment effect by household's equivalised income, Y_i . All regressions include household fixed effects, year-month effects and controls for major holidays. Standard errors are shown in parentheses below the coefficient estimates.

A.4 Variety and size definition

For estimation of our demand model we aggregate the more than 7000 alcohol UPCs (barcodes) purchased in 2011 into 32 types in order to make the demand system

tractable, following the approach in Griffith, O’Connell, and Smith (2019). We, as much as possible, only aggregate across UPCs that are of a similar alcohol type, quality and price. See Table A.4 for a list of types.

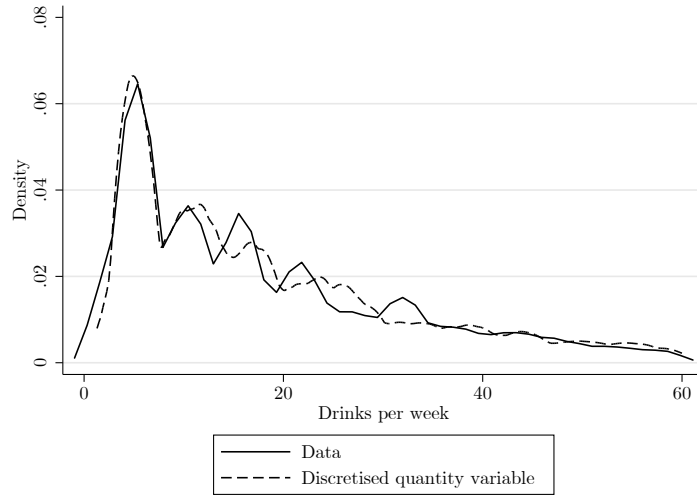
Table A.4: *Type definition and characteristics*

(1)	(2)	(3)	(4)	(5)	(6)
Type	Top brand and within-type share (%)	No. brands	Mean ABV	Market share (%)	No. sizes
<i>Beer</i>					
Premium beer; ABV < 5%	Newcastle Brown Ale (6.1)	386	4.4	1.8	3
Premium beer; ABV ≥ 5%	Old Speckled Hen (16.5)	238	5.5	2.1	3
Mid-range bottled beer	Budweiser Lager (19.6)	94	4.7	4.6	3
Mid-range canned beer; ABV < 4.5%	Carlsberg Lager (28.8)	17	3.9	5.8	3
Mid-range canned beer; ABV ≥ 4.5%	Stella Artois Lager (72.0)	15	5.0	2.7	3
Budget beer	John Smiths Bitter (23.6)	72	4.2	3.2	3
<i>Wine</i>					
Red wine	Tesco Wine (6.2)	435	12.6	18.4	4
White wine	Tesco Red Wine (7.8)	335	12.1	17.1	4
Rose wine	Echo Falls Wine (8.6)	63	11.5	4.2	2
Sparkling wine	Lambrini Sparkling Wine (8.4)	125	9.2	3.1	2
Champagne	Lanson Champagne (12.7)	42	11.8	0.8	1
Port	Dows Port (22.0)	23	19.8	0.7	1
Sherry	Harveys Bristol Cream (18.7)	25	16.8	1.2	1
Vermouth	Martini Extra Dry (11.8)	33	15.0	0.6	1
Other fortified wines	Tesco Fortified Wine (21.8)	37	14.6	0.9	1
<i>Spirits</i>					
Premium gin	Gordons Gin (59.6)	21	38.3	1.6	2
Budget gin	Tesco Gin (22.3)	15	38.3	1.3	2
Premium vodka	Smirnoff Red Vodka (39.0)	54	37.6	3.1	2
Budget vodka	Tesco Vodka (31.4)	17	37.5	1.8	2
Premium whiskey	Jack Daniels Bourbon/Rye (19.6)	80	40.5	2.1	2
Budget whiskey	Bells Scotch Whiskey (18.7)	56	40.0	8.1	2
Liqueurs; ABV <30%	Baileys (25.9)	203	18.4	3.1	2
Liqueurs; ABV ≥30%	Southern Comfort (27.2)	41	37.0	0.8	2
Brandy	Tesco Brandy (22.1)	55	37.3	2.4	2
Rum	Bacardi White Rum (29.1)	58	37.1	2.0	2
Pre-mixed spirits	Gordons Gin+Tonic (14.7)	43	6.1	0.2	1
Alcopops	Smirnoff Ice Vodka Mix (17.3)	147	4.8	0.8	1
<i>Cider</i>					
Apple cider, <5% ABV	Magners Original Cider (26.9)	52	4.4	1.6	3
Apple cider, 5-6% ABV	Strongbow Cider (63.1)	49	5.3	2.0	3
Apple cider, >6% ABV	Scrumpy Jack Cider (18.7)	71	7.0	0.8	2
Pear cider	Bulmers Pear Cider (24.2)	33	4.9	0.7	2
Fruit cider	Jacques Fruit Cider (21.4)	48	4.4	0.5	2

Notes: Replicated from Griffith, O’Connell, and Smith (2019). Column (1) shows the product definition. Column (2) lists the brand that constitutes the largest share of spending within each product; its within-product expenditure share is shown in parentheses. Column (3) lists the number of brands within each product. Column (4) shows the mean alcoholic strength (ABV) of each product. Column (5) shows the share of the alcohol market accounted for by each product. Column (6) shows the number of bins used to divide the quantity distribution.

Varieties are defined as combinations of types and the most commonly purchased quantities – see Table 4.1 in the main paper for a list of these. Figure A.2 shows that the distribution of drinks per adult per week across household-weeks in the data and constructed based on our discretization of the quantity distribution are very similar.

Figure A.2: *Drinks distribution with discretized size variable*



Notes: Replicated from Griffith, O’Connell, and Smith (2019). The solid line plots the distribution of drinks purchased per household-week calculated using the raw data. The dashed line plots the distribution of drinks purchased per household-week calculated using the discretized quantity variable.

B Additional details on the demand model

B.1 Identification and control function

We use a control function to help identify the slope of demand, as in Griffith, O’Connell, and Smith (2019). The instruments we use include changes in tax rates, exchange rates and factory gate prices.

Changes in tax rates applied to different alcohol products over our estimation period are shown in Table B.1.

Table B.1: *Tax changes during 2011*

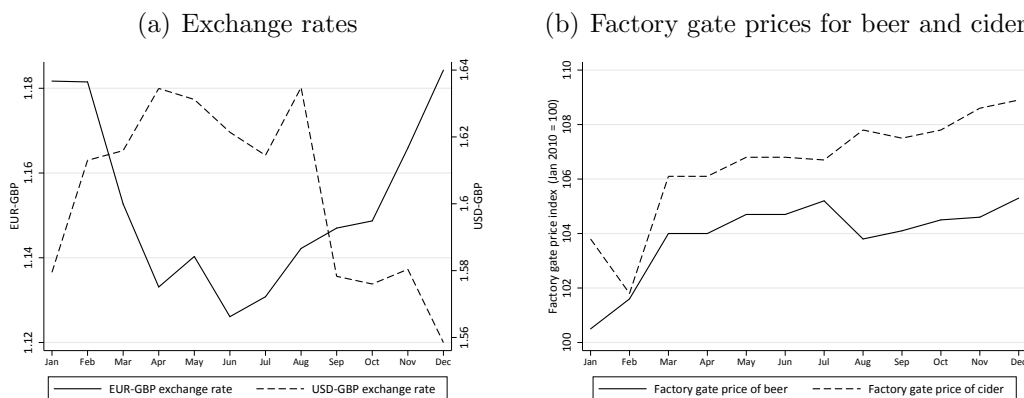
Segment	Applies to products:	Rate in Jan 2011:	Rate changes (month)
Beer	1.8-2.8% ABV	£17.32/litre ethanol	+1.25 (March); -9.28 (Oct)
	2.8-7.5% ABV	£17.32/litre ethanol	+1.25 (March)
	>7.5% ABV	£17.32/litre ethanol	+1.25 (March); +4.64 (Oct)
Wine	5.5-15% ABV (still)	£225.00/hectolitre product	+16.23 (March)
	15-22% ABV (still)	£299.97/hectolitre product	+21.64 (March)
	5.5-8.5% ABV (sparkling)	£217.83/hectolitre product	+15.72 (March)
	8.5-15% ABV (sparkling)	£288.20/hectolitre product	+20.79 (March)
Spirits	0-100% ABV	£23.80/litre ethanol	+1.72 (March)
Cider	1.2-7.5% ABV	£36.01/hectolitre product	-0.14 (March)
	7.5-8.5% ABV	£54.04/hectolitre product	-0.17 (March)

Notes: Alcohol duty is levied by alcohol segment and strength. The first column lists the combinations of segment and alcohol-by-volume (ABV) across which taxes vary. The second column shows the duty rate, and whether it was levied per unit of alcohol or per litre of product in January 2011. The third column lists the changes in duty rates, and, in parentheses, the month in which the change occurred. Tax rates from Her Majesty's Revenue and Customs.

There was considerable variation in the EUR-GBP and USD-GBP exchange rates, this is shown in Figure B.1(a). Movements in the exchange rate are likely to affect the prices of products differentially, depending on whether they are imported directly, or use imported inputs.

Figure B.1(b) shows that the factory gate prices for beer and cider changed differentially over 2011.

Figure B.1: *Exchange rates and factory gate prices, 2011*



Notes: Left hand panel plots the EUR-GBP and USD-GBP exchange rates over 2011 (data from Bank of England). Right hand panel plots the factory gate prices for beer and cider over 2011 (data from UK Office for National Statistics).

In the first stage regression our instruments are duty rates (interacted with options), exchange rates (interacted with options), beer and cider producer prices (interacted with beer and cider options), and oil prices (interacted with regions). The F-stat of the first stage is 36.40. In Table B.2 we show the tests of joint significance of different sets of variables.

Table B.2: *First stage: tests of instrument significance*

Set of instruments	F-statistic on joint significance
Alcohol duty rates (interacted with options)	3.25
Beer and cider producer prices (interacted with beer and cider options)	3.06
Exchange rates (interacted with options)	3.74
Oil prices (interacted with regions)	0.44
All instruments	36.40
Product effects included?	Yes
Flexible function of pack size included?	Yes
Type-time effects included?	Yes
Region effects included?	Yes

Notes: Table shows joint significance tests for sets of instruments for the first stage of the control function. First stage estimated at the option-region-month level; F-tests use heteroskedasticity-robust standard errors. Instruments are described in more detail in the text.

B.2 Estimates

Table B.3: *Estimated preference parameters*

Type of drinker:	Light			Moderate			Heavy		
Income group:	Low	Med	High	Low	Med	High	Low	Med	High
Panel A: Preferences for observable product characteristics									
<i>Means</i>									
Price	-0.360 (0.025)	-0.257 (0.024)	-0.174 (0.023)	-0.278 (0.021)	-0.227 (0.020)	-0.294 (0.019)	-0.419 (0.016)	-0.307 (0.017)	-0.179 (0.015)
Beer*Quantity of ethanol	0.158 (0.014)	0.110 (0.014)	0.080 (0.014)	0.149 (0.011)	0.123 (0.011)	0.177 (0.011)	0.239 (0.009)	0.189 (0.009)	0.109 (0.008)
Wine*Quantity of ethanol	0.021 (0.016)	-0.038 (0.015)	-0.087 (0.015)	0.040 (0.013)	0.023 (0.013)	0.051 (0.012)	0.176 (0.010)	0.102 (0.011)	0.045 (0.009)
Spirits*Quantity of ethanol	0.286 (0.028)	0.249 (0.027)	0.156 (0.026)	0.263 (0.023)	0.211 (0.023)	0.285 (0.022)	0.391 (0.018)	0.351 (0.019)	0.222 (0.017)
Cider*Quantity of ethanol	0.092 (0.015)	0.049 (0.015)	0.020 (0.014)	0.132 (0.012)	0.086 (0.011)	0.088 (0.012)	0.203 (0.009)	0.139 (0.010)	0.099 (0.010)
Beer*Quantity of ethanol ²	-0.001 (0.000)	-0.001 (0.000)	-0.001 (0.000)	-0.002 (0.000)	-0.002 (0.000)	-0.002 (0.000)	-0.002 (0.000)	-0.002 (0.000)	-0.001 (0.000)
Wine*Quantity of ethanol ²	0.001 (0.000)	0.002 (0.000)	0.002 (0.000)	0.001 (0.000)	0.001 (0.000)	0.001 (0.000)	0.000 (0.000)	0.001 (0.000)	0.001 (0.000)
Spirits*Quantity of ethanol ²	-0.003 (0.000)	-0.003 (0.000)	-0.002 (0.000)	-0.003 (0.000)	-0.002 (0.000)	-0.003 (0.000)	-0.004 (0.000)	-0.004 (0.000)	-0.002 (0.000)
Cider*Quantity of ethanol ²	-0.002 (0.000)	-0.001 (0.000)	-0.001 (0.000)	-0.002 (0.000)	-0.001 (0.000)	-0.001 (0.000)	-0.002 (0.000)	-0.002 (0.000)	-0.002 (0.000)
<i>Variances ×100</i>									
Price	1.031 (0.160)	0.520 (0.103)	0.117 (0.051)	1.641 (0.133)	1.270 (0.123)	0.101 (0.028)	1.299 (0.145)	0.009 (0.015)	0.679 (0.073)
Quantity of ethanol	0.295 (0.024)	0.286 (0.028)	0.339 (0.037)	0.567 (0.029)	0.214 (0.017)	0.097 (0.013)	0.421 (0.029)	0.178 (0.021)	0.349 (0.018)
Strength	0.242 (0.018)	0.319 (0.025)	0.175 (0.019)	0.397 (0.019)	0.287 (0.016)	0.322 (0.023)	0.464 (0.028)	0.361 (0.021)	0.350 (0.017)
<i>Covariances ×100</i>									
Price*Quantity of ethanol	-0.255 (0.053)	-0.215 (0.047)	-0.168 (0.046)	-0.694 (0.054)	-0.317 (0.044)	0.073 (0.006)	-0.540 (0.061)	-0.033 (0.029)	-0.357 (0.032)
Price*Alcohol strength	-0.026 (0.023)	-0.236 (0.034)	0.005 (0.010)	-0.072 (0.020)	-0.391 (0.028)	-0.162 (0.023)	-0.441 (0.037)	-0.004 (0.005)	-0.119 (0.014)
Quantity of ethanol*Alcohol strength	-0.122 (0.018)	-0.060 (0.023)	-0.136 (0.019)	-0.268 (0.017)	-0.052 (0.013)	-0.128 (0.015)	-0.046 (0.014)	-0.122 (0.012)	-0.132 (0.011)
Panel B: Preferences for unobserved product characteristics									
<i>Mean product effects for each segment</i>									
Beer	-5.812 (0.153)	-5.210 (0.130)	-5.819 (0.137)	-4.044 (0.111)	-4.500 (0.127)	-4.135 (0.151)	-3.754 (0.115)	-3.750 (0.114)	-3.148 (0.147)
Wine	-3.471 (0.178)	-3.204 (0.153)	-3.314 (0.151)	-3.076 (0.126)	-3.226 (0.132)	-2.238 (0.152)	-2.275 (0.125)	-2.282 (0.124)	-2.186 (0.151)
Spirits	-8.028 (0.296)	-8.233 (0.282)	-8.177 (0.289)	-6.459 (0.245)	-6.505 (0.260)	-6.833 (0.271)	-6.256 (0.210)	-6.799 (0.234)	-5.933 (0.242)
Cider	-5.366 (0.155)	-4.796 (0.136)	-4.998 (0.135)	-4.450 (0.118)	-4.102 (0.126)	-3.952 (0.142)	-3.496 (0.109)	-3.947 (0.121)	-3.570 (0.135)
<i>Variances</i>									
Beer	2.335 (0.126)	2.491 (0.142)	1.926 (0.128)	2.642 (0.117)	3.405 (0.195)	2.994 (0.249)	2.188 (0.103)	2.325 (0.122)	1.675 (0.093)
Wine	1.707 (0.099)	1.691 (0.098)	1.563 (0.092)	1.881 (0.090)	1.401 (0.067)	1.281 (0.068)	2.498 (0.109)	2.059 (0.100)	1.265 (0.070)
Spirits	0.401 (0.068)	0.940 (0.095)	0.992 (0.089)	0.667 (0.059)	1.645 (0.110)	0.869 (0.074)	0.475 (0.040)	0.443 (0.050)	0.910 (0.097)
Cider	4.602 (0.281)	2.665 (0.213)	2.207 (0.172)	3.034 (0.158)	4.401 (0.262)	3.053 (0.191)	3.985 (0.195)	3.148 (0.237)	2.466 (0.131)
Control function	0.242 (0.031)	0.170 (0.029)	0.075 (0.028)	0.135 (0.024)	0.157 (0.023)	0.272 (0.022)	0.372 (0.017)	0.260 (0.017)	0.185 (0.015)
Product effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Type-time effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Outside option-region effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

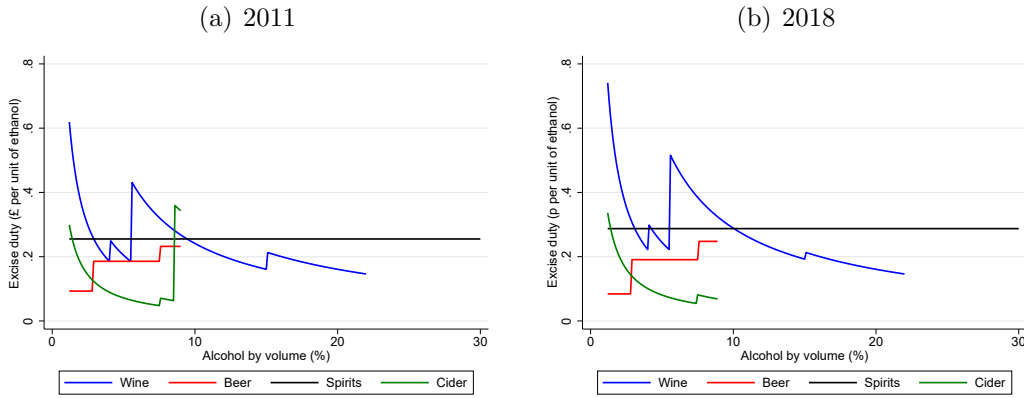
Notes: Light drinkers buy fewer than 7 units per adult per week on average (measured over a year), moderate drinkers between 7 and 14 units, and heavy drinkers more than 14 units. Panel A shows estimated parameters for the distribution of preferences over observable product characteristics, Panel B shows estimated parameters for the distribution of preferences over unobserved product characteristics. Standard errors are reported below the coefficients.

C Additional details on policy reforms

C.1 Alcohol duties

The UK implements a system of alcohol duties that vary across alcohol segment and by alcoholic strength; these are illustrated for 2011 and 2018 in Figure C.1. Wine and cider are taxed volumetrically i.e., per litre sold, with rates varying across different strength bands, which leads to the downward sloping kinked lines when the rates are expressed per unit of alcohol. Spirits and cider are taxed per unit of pure ethanol. Table B.1 summarizes the duty changes over the period that we estimate demand that we use to help identify the impact of price coefficient.

Figure C.1: *UK system of alcohol duties*



Notes: Figure shows the excise duty (expressed per unit) levied on different alcohol segments in 2011 and 2018.

C.2 Pass-through of taxes to consumer prices

In Section 5 we assume that taxes are fully passed through to consumer prices i.e., there is no over or undershifting. Here we present supporting empirical evidence.

Over the period 2010-12 there were several changes in the duty levied on different types of alcohol – these changes are summarized in Table B.1. We look at the effect of this on prices to assess whether pass-through is complete. As in Section 3, we let j denote alcohol product, and t denote transaction; ρ_{jt} denotes the transaction price of j on t . We let τ_{jt}^d denote the total amount of duty liable on product j on transaction t , and let τ_t^v denote the VAT rate applied to transaction t . Let $\tilde{\rho}_{jt}$ denote the tax-exclusive price:

$$\rho_{jt} = (1 + \tau_t^v)(\tilde{\rho}_{jt} + \tau_{jt}^d)$$

First, we estimate whether changes in duty lead to changes in the tax-exclusive price, $\tilde{\rho}_{jt}$, following an approach similar to that of Besley and Rosen (1999):

$$\ln \tilde{\rho}_{jt} = \beta \tau_{jt}^d + \mu_j + \gamma_{r(t)} + \xi_{m(t)} + \zeta_{y(t)} + \epsilon_{jt} \quad (\text{C.1})$$

where μ_j are product fixed effects, $\gamma_{r(t)}$ are retailer fixed effects, $\xi_{m(t)}$ and $\zeta_{y(t)}$ are month and year fixed effects. The coefficient of interest is β . If $\beta > 0$ then there is over-shifting of the tax (the tax-exclusive price increases when the tax rises), and $\beta < 0$ corresponds to undershifting. If we cannot reject $\beta = 0$ then we cannot reject complete pass-through of the tax to prices. Column (1) in Table C.1 shows that we cannot reject $\beta = 0$.

Second, we estimate an alternative specification that uses the fact that:

$$\frac{d\rho_{jt}}{d\tau_{jt}^d} = (1 + \tau_t^v) + \frac{d\tilde{\rho}_{jt}}{d\tau_{jt}^d}. \quad (\text{C.2})$$

If pass-through is complete, then $\frac{d\rho_{jt}}{d\tau_{jt}^d} = (1 + \tau_t^v)$. We therefore estimate:

$$\rho_{jt} = \alpha \tau_{jt}^d + \mu_j + \gamma_{r(t)} + \xi_{m(t)} + \zeta_{y(t)} + \epsilon_{jt}$$

and test whether α is significantly different from $1 + \tau_t^v = 1.2$. Column (2) in Table C.1 shows that we cannot reject $\alpha = 1 + \tau_t^v$.

Table C.1: *Tax pass-through*

	(1)	(2)
	Log tax-exclusive price	Price
Alcohol duty liable	-0.03 [-0.08,0.03]	1.21 [1.09,1.34]
Product effects	Yes	Yes
Month effects	Yes	Yes
Year effects	Yes	Yes
Retailer effects	Yes	Yes

Notes: Column (1) shows the estimated $\hat{\beta}$ from equation (C.1) and column (2) shows the estimated $\hat{\alpha}$ from equation (C.2). The square brackets show 95% confidence intervals; standard errors are heteroskedasticity robust and clustered at the product level.

C.3 Externality function

Table C.2: *Externality function calibration*

(1)	(2)	(3)		(4)	
Value of	% external costs	Marginal externality (£)			
b	from heavy drinkers	Median drinker	90th percentile	Ratio	
1.0	49.1	0.28	0.28	1.0	
1.3	57.6	0.26	0.42	1.6	
1.6	65.0	0.22	0.58	2.6	
1.9	71.3	0.17	0.74	4.3	
2.2	76.4	0.13	0.89	6.9	
2.5	80.6	0.09	1.02	11.2	
2.8	84.0	0.06	1.14	18.2	
3.1	86.7	0.04	1.23	29.5	
3.4	88.9	0.03	1.30	47.8	
3.7	90.6	0.02	1.35	77.5	
4.0	92.1	0.01	1.37	125.7	

Notes: Column (2) shows the share of external costs created by heavy drinkers (those that purchase more than 14 units per adult per week) under the existing system for different calibrated values of b . Columns (3) and (4) show the marginal externality of the median drinker and the drinker at the 90th percentile, and column (5) shows the ratio of (4) to (3).