

Time Inconsistency and Alcohol Sales Restrictions ^{*}

Marit Hinnosaar [†]

May 2016

Abstract

Restrictions on alcohol sales hours or days are commonly used tools in order to reduce alcohol consumption. However, a forward-looking consumer can buy in advance, and thereby mostly undo the impact of the restriction. I study whether time inconsistent consumer preferences can provide a justification for restrictions on alcohol sales time. I estimate a demand model, which allows a fraction of consumers to be time inconsistent, using scanner data of beer purchases and other shopping behavior. According to the estimation results, 16% of regular beer buyers, or only 3% of all consumers, behave as if they are time inconsistent. I find that in terms of consumer welfare, the sales restriction may be welfare improving, but is worse than increasing taxes.

Keywords: behavioral economics, time inconsistent preferences, consumer demand, alcohol, public health policy.

JEL Codes: D03, D12, I18, L51, L66.

^{*}The paper was a part of my dissertation at Northwestern University. For their guidance and support, I am grateful to my advisor Aviv Nevo, and my dissertation committee members Matthias Doepke, Rob Porter, and Shane Greenstein. For many useful comments I thank Tiago Pires, Ka Hei Tse, Chris Vickers, Mary Zaki, Jose Miguel Abito, Lorenzo Michelozzi, Mark Chicu, Arkadiusz Szydlowski, Jose Espin, Esteban Petruzzello, Andrea Pozzi, Marciano Siniscalchi, Todd Sarver, Ainhoa Aparicio, Michael Kummer, Ignacio Monzon, and Toomas Hinnosaar. I am grateful for the discussions and feedback at the seminars in Northwestern University, University of Copenhagen, CERGE-EI, New Economic School, Collegio Carlo Alberto, EIEF, University of Queensland, University of Oslo, Aalto University, the Berlin Colloquium in Behavioral Economics, University of Naples, the MOOD Workshop in Rome, IIOC in Boston, Barcelona GSE Summer Forum, EEA meeting in Gothenburg, and EARIE in Evora.

[†]Collegio Carlo Alberto, E-mail: marit.hinnosaar@gmail.com

1 Introduction

In many countries, public health policy would like to limit alcohol consumption. One commonly used measure is to restrict the days when alcohol is available in stores and health authorities argue that these restrictions are good tools to reduce alcohol consumption. However, it is not clear why the restrictions work. A forward-looking consumer should be able to buy in advance, and thereby mostly undo the impact of the restriction. In this paper, I consider two explanations as to why the restrictions might work. First, the restrictions make it more inconvenient to buy alcohol. Second, some consumers might have time inconsistent preferences.¹ I estimate a demand model to evaluate these explanations and study whether time inconsistent preferences can provide a justification for restrictions on alcohol sales time.

Consider a time inconsistent consumer. This is a person with a self-control problem. He enjoys a drink today, while promising himself that he will not be drinking tomorrow. But when tomorrow arrives, he repeats his behavior. He decides to drink again and hopes that he will drink less in the future.² For this person the regulation provides a commitment device that helps him to limit his consumption.

In this paper, I estimate the fraction of consumers who behave as if they are time inconsistent and study the implications for policy. Specifically, I ask whether sales restrictions can improve consumer welfare and be preferable to taxes. To answer the questions, I first construct a model of alcohol purchasing, consuming, and storing. The model allows consumers to be time inconsistent and includes time consistency as a special case. Then I estimate the model, including consumers' time inconsistent discount factor, using scanner data of beer purchases and other shopping behavior. My data set is from regions with no sales restrictions and has information about the daily purchases of more than 5000 regular beer buyers over eleven years. I use these estimates to evaluate the policy.

I identify time inconsistent behavior from dynamic purchasing patterns. Consider again the time inconsistent consumer. When he is in the store on Saturday and plans not to return the next day, then there are two important aspects of his behavior. First, on average he buys less for tomorrow than for current consumption, since he believes he wants to decrease his consumption in the future. Second, since he underestimates how much he wants to consume

¹The literature provides evidence that some people might behave as if they are time inconsistent (Frederick, Loewenstein, and O'Donoghue, 2002) and alcohol consumption is a natural setting where we would expect to see evidence of time inconsistency.

²We say that a consumer who prefers a drink today to 5 dollars tomorrow, but reverses his decision when both offers are further postponed by a day, exhibits time inconsistent preferences regarding alcohol. This can be thought of as a special case of the definition that involves preference reversal over time.

tomorrow, he sometimes returns to the store when he did not plan to. To summarize this, we expect the time inconsistent consumer to purchase more for current consumption than for future consumption, and to return to the store more often for additional purchases. Of course, in the data we only observe total purchases, and cannot separate inventory from current consumption. However, I will show that we can partially distinguish these two, assuming optimizing behavior and variation in the shopping costs. This is so, since when the shopping cost in the future is high, the consumer is more likely to purchase inventory than in the periods when it is costless to purchase additional quantity later.

In the model that I will describe below, the motivation for storing inventory comes from shopping costs. That is, the consumer purchases for future consumption in order to avoid a shopping trip every time he would like to consume something. This is different from most of the storable good demand models (e.g. Hendel and Nevo, 2006). In those, the consumer wants to purchase for future consumption when the good is on sale and he expects the price to increase. The reason for not considering price changes in this paper is that when I turn to the data, I look at purchasing decisions in the short-term horizon, over the weekend, and prices usually remain constant during this period. In the short-term horizon, it is reasonable to expect that the main motivation for storing inventory comes from shopping costs. The role of shopping costs in the model is similar to the role of sales in Hendel and Nevo (2013). Namely, if the future shopping costs are high, the consumer may want to store inventory in order to avoid paying the high shopping cost.

Identification of time inconsistent preferences is based on two key assumptions. First, shopping costs are partially observable to the researcher. Specifically, I assume that when a consumer purchased a large shopping basket of goods other than beer, he was visiting the store for reasons other than beer, and therefore his shopping cost for beer had to be relatively low. Second, I assume that shopping costs are uncorrelated with the daily preference shocks that determine the wish to consume beer on that day. That is, after controlling for observable characteristics that might affect beer demand (like day of the week and seasonal effects), I assume that shopping costs and preference shocks are not correlated.

According to the estimation results, 16% of regular beer buyers, a group which makes up 20% of all consumers, behave as if they are time inconsistent, having a time inconsistent discount factor equal to 0.47. But since these consumers consume more beer, their share in the beer consumption of regular beer buyers is much larger, namely 52%. This drives the welfare implications.

Using the estimated parameters of the model, I run two counterfactual policy experiments.

Specifically, I look at what happens with consumer welfare when a Sunday sales restriction is introduced or taxes are increased. I consider the case where tax revenue is redistributed to the consumers as a lump-sum payment. When comparing the sales restriction and taxes, my goal is to compare two policies that can achieve the same decrease in consumption. I find that the sales restriction may increase consumer welfare. However, the sales restriction is worse than taxes. Note that the sales restriction has two effects. First, it increases costs by making it more inconvenient to buy. Second, for the time inconsistent consumers, it provides a commitment device. In comparison, taxes are more flexible in allowing consumers to choose when to decrease consumption. Since time consistent consumers do not need the commitment, they don't benefit from the restriction and would always prefer taxes. This is in contrast to the time inconsistent consumers. The welfare impact of the sales restriction depends on whether the commitment effect of time inconsistent consumers is large enough to compensate for the increased inconvenience.

I should note that the estimates are based on a specific sample of consumers. First, the main sample is restricted to the consumers who regularly buy beer. However, in the counterfactual analysis, I present back-of-the-envelope calculations that extrapolate the results to the full sample. Second, the consumers are from regions where the beer regulation is liberal. The estimates and results from the counterfactual policy experiment might not generalize to a population in other areas. Finally, note that the analysis here is based only on beer purchases in stores. When designing a policy, substitutes like other types of alcohol and consumption in bars should be taken into account.

This paper is related to the literature that studies alcohol markets and policy, and specifically, the usefulness of sales restrictions in decreasing alcohol consumption. In the literature, the empirical evidence regarding the impact of alcohol off-premise sales restrictions is mixed. The datasets used to measure the impact usually include some policy changes. But even in the case of a policy change, it is difficult to identify the impact, since relaxing the sales restriction is often accompanied by changes in other restrictions and taxes. Moreover, the policy change itself is usually a political decision, taking place due to changes in voter preferences. Perhaps the best dataset to study the impact of sales restrictions comes from a controlled policy experiment in Sweden. In 2000, the Saturday sales ban was lifted in randomly chosen counties, while in the control counties the restriction remained in place for the following year and a half. Using data from this experiment, Norström and Skog (2003, 2005) find that total alcohol sales increased by 3%, and beer sales increased even more, by 8%. There are many other papers that have looked at changes in alcohol consumption associated

with policy changes. But as said above, the results have been mixed, perhaps because the environment of the policy changes has been less controlled. For example, Stehr (2007) finds that after the repeal of Sunday sales bans in the U.S., liquor but not beer sales increased. On the other hand, Bernheim, Meer, and Novarro (2016) also look at policy changes in the U.S., and Carpenter and Eisenberg (2009) look at policy changes in Canada, and they both find no effect on overall alcohol consumption.^{3,4} In my dataset there is no policy change. But, based on my estimates the sales restriction leads to a predicted decrease in beer consumption by 10%. This implies that the impact on total consumption is in line but somewhat larger than in the findings from the policy experiment in Sweden.

At a more general level, the paper is related to the literature that identifies time inconsistent preferences and then measures their empirical importance outside the lab. This literature has taken three different paths. First, there are papers that combine lab and field data. In those, the discount rate is learned using lab experiments, and then this is combined with real world data about the individual. Examples include credit card usage (Meier and Sprenger, 2010), technology adoption (Tarozzi and Mahajan, 2011), and studying how present bias is related to income (Tanaka, Camerer, and Nguyen, 2010). Second, there are papers where the identification is based on policy experiments or discontinuities in policies or contracts. For example, Shapiro (2005), Mastrobuoni and Weinberg (2009), and Hastings and Washington (2010) study food consumption before and after the arrival of social security paychecks, while DellaVigna and Malmendier (2006) study gym membership. The third branch of papers use dynamic models to identify time inconsistent preferences. The settings include lifetime consumption-saving decisions (Laibson, Repetto, and Tobacman, 2007), unemployment and labor supply (Paserman, 2008; Fang and Silverman, 2009), and mammography decisions (Fang and Wang, 2015). In this paper, I follow the last approach and identify a time inconsistent discount rate from dynamic choices. My estimated time inconsistent discount factor equals 0.47. Note that most studies pool all consumers together and estimate a common time inconsistent discount factor, while according to my estimates, the majority of consumers have a time inconsistent discount factor equal to one, and only 3% (in the full sample) discount at

³In addition to studying the direct effect on consumption, many papers have analyzed the impact on alcohol related harms like drunk driving, traffic accidents, violence, and health conditions related to alcohol. For an overview see Popova, Giesbrecht, Bekmuradov, and Patra (2009).

⁴There is some evidence of the heterogeneous impact on individuals. For example, Gruber and Hungerman (2008) looked at the impact of repealing Sunday restrictions on regular commerce, and found that it increased heavy drinking (consuming 6 or more drinks in one sitting in the past month) only among the individuals who before the repeal regularly attended church. Marcus and Siedler (2015) found that a late night alcohol sales ban reduced alcohol-related hospitalizations only among young people and not among older individuals.

the higher rate.

This paper contributes to the literature on optimal regulation of "unhealthy" goods in the case of time inconsistent consumers. In theoretical models, Gruber and Köszegi (2001) and O'Donoghue and Rabin (2006) find the optimal "sin" taxes in the case of time inconsistent consumers who at the time of consumption underweigh the future costs. They find that the optimal tax rate is higher than with time consistent consumers.⁵ All these papers look only at taxes and do not consider sales restrictions. The impact of sales restrictions on consumers who are not time consistent has been analyzed in a theoretical model by Beshears, Choi, Laibson, and Madrian (2006).⁶ They show that when restricting sales is costless, sales restrictions are optimal.⁷ They mention that the regulation comes with social costs, including increased inconvenience, which they do not include in their model, but due to those costs the welfare impact of the regulation becomes an empirical question. In my model, heterogeneous shopping costs capture this inconvenience, and the sales restriction is not necessarily the optimal policy, since in some situations it makes it more costly to purchase alcohol.

The rest of the paper is organized as follows. In Section 2, I discuss alcohol regulation in the U.S., describe my data and look at the descriptive empirical evidence. Section 3 presents the model, Section 4 discusses the estimation and Section 5 presents the estimation results. In Section 6, I present the counterfactual policy analysis to evaluate the impact of the regulation. First, I study the impact of the sales restriction on quantity and welfare, and then compare this to the impact of taxes. Section 7 concludes.

⁵Papers with similar findings include Gruber and Köszegi (2004) and Haavio and Kotakorpi (2011).

⁶Bernheim and Rangel (2004) also consider a somewhat similar setting when studying addiction and cue-triggered decision making, which is not hyperbolic discounting but still deviates from time consistent preferences. They analyze supply disruptions (which include sales restrictions) and find that this could be a beneficial policy measure.

⁷In their model, restricting sales at suitable times provides a costless commitment device. That is, the sales restriction does not change the long-term self's payoff and eliminates the choices of the short-term self. Thereby, the restriction provides a way for the long-term self to maximize his payoff.

2 Institutional Background, Data, and Descriptive Analysis

2.1 Institutional Background

In the U.S., as well as in other countries, the public health agenda includes objectives to decrease alcohol consumption.⁸ One of the policy measures suggested is to restrict the hours and days when retail stores can sell alcohol. For example, in 2010, 193 member states of the World Health Organization adopted a resolution on a global strategy to reduce the harmful use of alcohol, and one of the policy measures suggested was regulating hours and days of retail sales.⁹

Restrictions on the hours and days when alcohol is allowed to be sold in stores are common in Europe and the U.S. In the U.S., several states and many local governments do not allow alcohol sales on Sundays and evenings. The restrictions in the U.S. go back to the Prohibition era, when alcohol sales were not allowed at all, and the restrictions have remained in place partly for religious reasons. During the past decade, several states have repealed the sales restrictions. While in 2001, 15 states did not allow off-premise beer sales on Sunday, by 2011 five of those states had repealed the ban. It is argued that the main reason for lifting the restriction has been the hope to increase tax revenues. But, as noted above, public health agendas stress that restrictions on sales time are important policy tools in reducing alcohol consumption. Therefore, in the popular press, public health officials have voiced concern that repealing the sales bans has negative health consequences.

The dataset used in this paper is from the following U.S. states, which all have liberal alcohol laws and allow alcohol sales on Sundays: Arizona, California, Florida, Illinois, Louisiana, Missouri, Nebraska, Nevada, South Dakota, and Wisconsin. In these states, there is no state control or monopoly of any type of alcohol, there is no Sunday off-premise sales restriction of any type of alcohol, and at least beer and wine are allowed to be sold in grocery stores. In the mainland U.S., only the ten states listed above fulfill all the three criteria in 2004 to 2014, which is the time period in my dataset. I restrict attention to the states with

⁸U.S. Department of Health and Human Services. Healthy People 2020. <https://www.healthypeople.gov/2020/topics-objectives/topic/substance-abuse/objectives/>, accessed April 23, 2016; World Health Organization. 2014. "Global status report on alcohol and health." http://www.who.int/substance_abuse/publications/global_alcohol_report/en/, accessed April 23, 2016.

⁹Resolution of the Sixty-third World Health Assembly. May 2010. WHA63.13 Global strategy to reduce the harmful use of alcohol, page 14. http://www.who.int/substance_abuse/alcstratenglishfinal.pdf, accessed April 23, 2016.

no Sunday sales restriction in order to see consumer shopping behavior on the weekends without the restriction. I require that grocery stores can sell alcohol, to look at the states where consumers can combine their grocery shopping and alcohol purchases, and don't have to make extra trips to liquor stores.

Alcohol consumption in the states above is somewhat higher than the U.S. average.

- Alcohol consumption in these states was 2.32 – 3.31 gallons of ethanol per capita per year compared to 2.34 as the U.S. average in 2013.¹⁰
- Beer consumption in these states was 18.4 – 27.5 gallons of beer per capita per year compared to 20.7 as the U.S. average in 2010.¹¹
- According to the Behavioral Risk Factor Surveillance System 2014 survey, the percentage of adults who had had at least one drink in the past 30 days in these states was 45.5 – 63.1% compared to 53.3% in the median state in the U.S in 2014.
- According to the same source, prevalence of heavy drinking in these states was 5.0 – 7.2% compared to 5.9% in the median state in the U.S in 2014.¹²

2.2 Data

I use consumer level scanner data from the Nielsen Homescan Consumer Panel.¹³ For each day and each consumer, the dataset includes the quantity and price of beer purchased, the time of all shopping trips, including the ones when beer was purchased, purchases of other goods, and consumer characteristics. The data is collected using in-home scanners that the consumers themselves use to scan all their grocery purchases intended for personal in-home use and purchased from any outlet. The dataset covers the time period from 2004 to 2014 and includes data from the 10 states that have liberal alcohol laws as described in Subsection 2.1: Arizona, California, Florida, Illinois, Louisiana, Missouri, Nebraska, Nevada, South Dakota, and Wisconsin.

¹⁰National Institute on Alcohol Abuse and Alcoholism. 2015. "Apparent Per Capita Alcohol Consumption: National, State, and Regional Trends, 1977 – 2013."

¹¹Beer Institute. 2011. "Brewers Almanac, 2011."

¹²The Centers for Disease Control and Prevention define heavy drinking as more than two drinks per day for men and more than one drink per day for women.

¹³The dataset is from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. Information on availability and access to the data is available at <http://research.chicagobooth.edu/nielsen>.

For most of the empirical analysis, I restrict attention to consumers who purchase beer often. A household in a given year is included in the main sample of frequent consumers if it purchased beer in at least 10 weeks per year and purchased beer in at least 5 weekends during the observable years. These households can be different over the years. If a household satisfies these criteria one year, but not in the others, then it is included in the sample for only that year. The exact procedures undertaken to clean the raw data are described in Appendix A. From the total of 27,419 households, 5385 households satisfy these criteria for at least one year.

The 20% of households who form the sample of frequent consumers used in the main empirical analysis account for about 71% of total beer sales. Table 1 displays descriptive statistics of both types of households. In terms of demographic characteristics, the frequent consumers are similar to the rest. However, the median household in the main sample spends more than ten times more on beer than the median household that does not purchase beer regularly and is therefore excluded from the main sample.

The median household in the main sample of frequent consumers has two members and a total income in the range of \$50,000 to \$59,999. The median household head is 50 to 54 years old. 18% of the households include a male who does not work full time. 38% of the households include a full-time working female. There are children in 26% of the households.

The median household in the main sample shops on average 1.8 days per week. Figure 1 displays the histogram of households by their average number of shopping days per week. The households in the sample buy beer rather often: the median households buys beer on average 0.4 days per week, which means about once in two weeks. Figure 2 displays the histogram of households by their average number of beer shopping days per week. Households buy beer throughout the week, however, the largest share of beer expenditure is on the weekend. Figure 3 presents the distribution of beer expenditure over the days of the week.

The quantity of beer bought in one shopping trip (conditional that beer is purchased) is quite large. The quantity is a six-pack (6 12-ounce bottles/cans) or less in only 17% of trips. On the median shopping trip when beer is bought, the quantity is 18 12-ounce bottles/cans. In 28% of trips, the quantity is more than a case (24 12-ounce bottles/cans) of beer. Figure 4 presents the distribution of quantities bought in one shopping trip. The figure also shows that the distribution of beer quantities bought is similar over weekdays and the weekend.

Table 2 displays descriptive statistics of the households' shopping trips. On a median weekend shopping trip, a household spends \$50. On a median weekend trip with a beer purchase, it buys 18 12 ounce bottles/cans. The average price of 24 12-ounce bottles/cans of

beer is about \$17.¹⁴ A household on a median weekend trip with a beer purchase spends \$13 on beer.

2.3 Descriptive Analysis

The general intuition is that the time inconsistent consumer buys more often. Figure 5 presents the percentage of consumers returning to buy beer given that they bought beer the previous day. 10% of households that bought beer on Saturday return to the store to buy beer on Sunday. In case of single-person households, the returning rate is similar. Even the percentage of households returning to the same type of store (grocery stores, liquor stores, convenience stores, etc) is still above 6%.

To see which consumers behave this way, I examine the frequency of buying beer regressed on various household characteristics. In regressions in Table 3, the dependent variable is an indicator for purchasing beer conditional on a beer purchase the previous day. Higher income and higher education decrease the likelihood of consecutive beer purchases. Households without a male are less likely to buy beer frequently. However, when we restrict attention to single person households, only gender seems to matter. Single women are much less likely to buy beer on consecutive days than single men. This is especially true if the women are working full time.

How does returning to buy beer depend on previous quantity bought? Column 1 in Table 4 presents the same regression as above, conditioning on the quantity bought the previous day. The smaller the quantity bought, the more likely the household is to buy again the next day. Which households buy small quantities? Columns 2 and 3 present estimates where the dependent variable is an indicator for buying a six-pack or less, conditional on buying beer. Households that tend to buy smaller quantities are more likely to be single, have a female working full time, and be more highly educated. How does the choice of quantity today depend on previous purchase? Column 2 includes an indicator whether beer was purchased the previous day. Given that the household purchases beer today, it is more likely to buy a small quantity when beer was bought the previous day. This hints towards a pattern where households that buy frequently are more likely to buy small quantities. Finally, Column 3 restricts attention to days when the household buys beer on two consecutive days. The estimates show that households that buy a small quantity on the first day are more likely to buy a small quantity on the second day.

¹⁴The price is calculated as a weighted average of the prices of all beer sales in the state in each week; prices are weighted by sales.

3 Model

In this section, I present the model, and then derive the equations which describe consumer behavior. These equations will be used in the next section for the structural estimation.

Setup: There are two goods, alcohol and a composite good, and three time periods. Timing is the following. A consumer buys and consumes alcohol on Saturday and Sunday, and consumes the composite good in the following week. That is, in periods $t = 1, 2$ he decides how much alcohol to purchase x_t at price p and how much to consume q_t . In the third period, he consumes the composite good. Price p is known to the consumer and remains constant during these periods.

Utility is quasilinear in the composite good (money). The payoff from alcohol consumption has two components: an instant payoff and a future cost. The instant payoff is logarithmic in the amount of alcohol consumed $\log(q_t + \eta_t)$, and is affected by a random preference shock η_t that affects the marginal utility of alcohol consumption and decreases the wish to consume. I assume that the preference shocks η_1, η_2 are independently and identically distributed on $[0, \infty)$, according to a distribution with cumulative density functions F_1, F_2 . The future cost is linear in the amount of alcohol consumed and the consumer bears that cost the following week. The future cost consists of a health cost $h \cdot q_t$ and alcohol expenditure $p \cdot x_t$.

Shopping trips are sometimes costly, with a fixed shopping cost per trip. The shopping cost can take one of two values $C_t \in \{0, C\}$, where $C > 0$. The value of the shopping cost does not depend on the quantity purchased. Storing is costless, but if indifferent, then the consumer chooses not to store any inventory.¹⁵

The uncertainty regarding consumption shocks (η_1, η_2) and shopping costs (C_1, C_2) realizes before period one. That is, from the consumer's perspective there is no uncertainty in the model.

Discounting: I consider discounting in the quasi-hyperbolic discounting framework (Laibson, 1997), sometimes called $\beta - \delta$ model, where β is the time inconsistent discount factor and δ is the time consistent exponential discount factor. I assume that $\delta = 1$. The results should not be sensitive to the assumption, because we are analyzing only one-day-ahead discounting.

There are two types of consumers who differ by their time inconsistent discount factor β . A Time Inconsistent (TIC) consumer has $\beta < 1$, and hence discounts his current payoff less

¹⁵The indifference assumption is made to simplify the exposition in this section. In the empirical analysis, I don't rely on the assumption, as discussed in Section 4.1.

than his future payoffs. I assume that the time inconsistent consumer is unaware of his future time inconsistency problem, meaning that he incorrectly believes that he will become time consistent in the future. In the literature, this is called the naive time inconsistent agent. A Time Consistent (TC) consumer has the time inconsistent discount factor equal to one, and hence weights all current and future payoffs equally.

Consumer optimization problem: At period one, each consumer chooses his optimal alcohol consumption q_1 , alcohol purchases x_1 , and alcohol inventory stored for the second period i_2 , given price p , current and future shopping costs, and current and next period alcohol consumption shocks. The consumer has the following payoff at period one:

$$U_1(p, \beta) = \log(q_1 + \eta_1) - C_1 \cdot 1[x_1 > 0] - \beta(hq_1 + \alpha px_1) + \beta [\log(q_2 + \eta_2) - C_2 \cdot 1[x_2 > 0] - (hq_2 + \alpha px_2)] \quad (1)$$

where η_1, η_2 are alcohol consumption shocks, C_1, C_2 are shopping costs, h is the health cost per unit of alcohol consumed, and α is the marginal utility of income. Payoff function U_1 also depends on the current and future consumption shocks and shopping costs, and the parameters of the model, but I omit those to shorten the notation. The consumer's payoff at period two equals:

$$U_2(p, i_2, \beta) = \log(q_2 + \eta_2) - C_2 \cdot 1[x_2 > 0] - \beta(hq_2 + \alpha px_2) \quad (2)$$

Note that the time inconsistent consumer discounts the health cost and alcohol expenditure, but not the instant payoff that he receives from alcohol consumption. This is based on the assumption that both the health consequences and consumption of the composite good take place in the third period. The assumption is meant to capture the idea that at the moment of purchasing alcohol, one compares the cost to the opportunity to buy alternative goods in the future.¹⁶

¹⁶Note that the time inconsistent discounting modeled in this way is similar to the temptation goods in Banerjee and Mullainathan (2010).

The optimization problem in period one is:

$$\begin{aligned} \max_{q_1, q_2, x_1, x_2, i_2} \quad & U_1(p, \beta) \\ \text{s.t.} \quad & q_1 + i_2 \leq x_1 \\ & q_2 \leq x_2 + i_2 \end{aligned} \tag{3}$$

and in period two is:

$$\begin{aligned} \max_{q_2, x_2} \quad & U_2(p, i_2, \beta) \\ \text{s.t.} \quad & q_2 \leq x_2 + i_2 \end{aligned} \tag{4}$$

The shopping cost pairs (either zero or $C > 0$) define four possible states. In short, the states are denoted by 00, C0, 0C, and CC. When the future shopping cost equals zero, then the consumer does not need to buy inventory and the optimization problem can be solved in each period separately. I call the corresponding demand function single period demand. Let me first explain the general intuition of the single period demand and then present the formula. If it is optimal for the consumer to go shopping, then for current consumption he purchases the quantity equal to $\max\{0, \frac{1}{\beta(h+\alpha p)} - \eta_t\}$, which is derived directly from the first order condition. Note that in order to purchase, the benefits from consumption must be large enough to cover the shopping cost. This is not restrictive when shopping cost equals zero. But, when shopping cost is higher, then demand is lower due to the high shopping cost. When the consumer purchases only for current consumption and has no inventory, then it is optimal for him to go shopping up to the following threshold for the consumption shock: $\eta_t \leq \frac{-W(-e^{-1-C_t})}{\beta(h+\alpha p)}$, where W is a Lambert W function (also called product log function). For compactness, let me denote $-W(-e^{-1-C_t}) = \Lambda(C_t)$. The function $\Lambda : [0, \infty) \rightarrow [0, 1]$ is such that $\Lambda^{-1}(x) = x - \log(x) - 1$. Note that Λ is decreasing and $\Lambda(0) = 1$. This implies that when shopping cost equals zero, it is optimal for the consumer to go shopping up to the point where he would like to buy a positive quantity, that is, up to the consumption shock threshold: $\eta_t \leq \frac{1}{\beta(h+\alpha p)}$. But, when shopping cost is higher, then the threshold will be lower. I denote the single period demand (expected quantity bought without inventory) given the current period shopping cost $C_t \in \{0, C\}$ by Q . The single period demand equals:

$$Q^\beta(C_t) = \int_0^{\frac{\Lambda(C_t)}{\beta(h+\alpha p)}} \left(\frac{1}{\beta(h+\alpha p)} - \eta \right) dF_t(\eta) \tag{5}$$

where for the time consistent consumer β equals one. Note that the single period demand is decreasing in shopping cost C_t and discount factor β . In particular, the single period demand of the time consistent consumer $Q^1(C_t)$ is lower than the single period demand of the time inconsistent consumer $Q^\beta(C_t)$.

A consumer stores inventory only when shopping is costly in the next period, that is, in states 0C and CC. First, consider the state where shopping cost equals zero in period one and C in the second period. It is useful to compare this state with the one where shopping costs are zero in both periods. The time consistent consumer buys the same total amount in both states

$$E(x_1^{TC}) = \begin{cases} Q^1(0) & \text{in } 00 \\ 2 \cdot Q^1(0) & \text{in } 0C \end{cases}, \quad E(x_2^{TC}) = \begin{cases} Q^1(0) & \text{in } 00 \\ 0 & \text{in } 0C \end{cases}$$

where $Q^1(0)$ is his single period demand with zero shopping cost. The time inconsistent consumer buys less when he buys for the next period compared to when he buys for the current period. If the discount factor is high enough, $\beta > \Lambda(C)$, then:

$$E(x_1^{TIC}) = \begin{cases} Q^\beta(0) & \text{in } 00 \\ Q^\beta(0) + Q^1(0) & \text{in } 0C \end{cases}, \quad E(x_2^{TIC}) = \begin{cases} Q^\beta(0) & \text{in } 00 \\ 0 & \text{in } 0C \end{cases}$$

where his single period demand with zero shopping cost is $Q^\beta(0)$ and his inventory is $Q^1(0)$. If the discount factor is low, the time inconsistent consumer has positive demand also in period two in state 0C.

Let's summarize the consumer behavior in state 0C. For the time consistent consumer, shopping cost determines only the time of the shopping trips, but it has no impact on the total quantity. When it is costly to shop in the second period, but not in the first, the consumer just chooses his inventory equal to his second period optimal consumption. In the case of the time inconsistent consumer, shopping cost decreases his total quantity bought and he may even return to the store in the second period.

Finally, consider the state where shopping costs are high in both periods. In this state, the shopping costs decrease consumption of both types of consumers. The time consistent consumer correctly predicts his next period preferences and behavior, but high shopping cost sometimes prevents him from purchasing. The time inconsistent consumer incorrectly predicts his next period preferences, and therefore may go to the store twice over the weekend. For details see Appendix B.

Sales restriction: Let us consider what would happen if alcohol sales were prohibited on Sundays. In the model, the restriction is equivalent to increasing the second period shopping cost. I introduce a total ban on second period sales, so that the second period shopping cost approaches infinity. I assume that the distribution of first period shopping costs remains unchanged. The assumption is restrictive, implying that in response to the alcohol sales restriction, consumers don't change their overall shopping behavior, that is, the time when they buy other goods. The assumption is more suitable for consumers for whom alcohol is not that important in their weekly shopping basket and who have other constraints that determine their weekly shopping schedule.

For the time consistent consumer, restricting sales on Sunday has an impact only due to shopping costs. That is, only when shopping in the first period is costly does the sales restriction reduce his demand. For the time inconsistent consumer, the sales restriction reduces demand for two reasons. The first reason is the shopping costs, and this is the same as for the time consistent consumer. Second, when the time inconsistent consumer has to buy in advance, then he buys less.

On consumer welfare the sales restriction has two opposing effects.¹⁷ First, it has a negative effect on consumer welfare, since it increases costs by making it more inconvenient to buy. Second, it has a positive effect, but only on time inconsistent consumers. For them it provides a device that helps them to decrease consumption and too frequent shopping trips. Without the sales restriction the consumption of the time inconsistent consumer was larger than his long-term self's welfare maximizing quantity. Let us summarize the welfare effects on these two types of consumers. Due to the shopping costs, the sales restriction always decreases the welfare of time consistent consumers. For time inconsistent consumers, the restriction can increase their welfare. It depends on the parameter values whether the commitment effect for time inconsistent consumers is large enough to compensate for their welfare loss from the increased inconvenience. This implies that, altogether, the sales restriction can increase total consumer welfare by helping time inconsistent consumers to make better decisions. Their welfare increase comes at the expense of time consistent consumers who will be worse off. Hence, total consumer welfare increases if first, for time inconsistent consumers the welfare gain from decreased consumption is larger than the welfare loss from higher shopping costs, and second, the share of time inconsistent consumers is large enough.

¹⁷I use the welfare criterion of the long-term self, which does not depend on the time inconsistent discount factor. For a detailed discussion of the alternative welfare criteria, see Section 6.1.

Discussion of assumptions about timing and discounting: In the model, I make the assumption that is common in the literature that the time inconsistent consumer discounts all his future cost and benefits by the same discount factor. The time inconsistent discount factor has two roles in the model. First, the time inconsistent consumer undervalues the costs associated with his current behavior. This implies that for him the relevant cost of consumption is essentially lower than for time consistent consumer, which implies that his demand will be higher. Second, the time inconsistent consumer also undervalues future consumption. This implies that when buying inventory he buys less than his next period self would have liked and therefore is likely to return to the store in the next period.

The dual role of the discount factor has implications for both the identification of time inconsistent consumers and welfare analysis. How to identify time inconsistent consumers? They: (a) are more likely to return to the store and buy beer in periods when shopping is costly, and (b) buy more beer, because for them the relevant price is essentially lower. In terms of their long term welfare, time inconsistent consumers consume too much and buy beer too frequently. The sales restriction reduces their consumption and too frequent shopping trips, which increases their welfare.

Let's consider an alternative assumption about discounting. Suppose the time inconsistent consumer does not discount the costs associated with his current consumption and purchases. Then his current period demand would equal that of the time consistent consumer. However, he would buy less inventory than the time inconsistent consumer and less than he himself would want to have in the next period (because the benefits related to future consumption will realize in the future and are discounted, while the cost of purchasing inventory realizes now) and this leads him to return to the store more often. So the model with the alternative discounting leads to the same results in terms of the time inconsistent consumer returning to the store more often, but to a different result in terms of consumption – the time inconsistent consumer would on average buy less than the time consistent consumer.

How do the two assumptions about discounting compare in terms of the welfare implications of the sales restriction? With the alternative assumption, in terms of their long term welfare, time inconsistent consumers don't consume enough beer, however, they buy beer too frequently. The sales restriction would have two opposing effects on time inconsistent consumers: it would lead to even lower consumption which decreases their welfare, and it reduces the number of costly shopping trips which increases their welfare. Hence, with this alternative assumption about discounting, the sales restriction is likely to be worse than with the standard assumption that I maintain in the paper.

Sophisticated time inconsistent consumer: Although the paper concentrates on the naive time inconsistent consumer, let me briefly discuss the behavior of the sophisticated time inconsistent consumer. The sophisticated consumer discounts the future in the same way as the naive consumer, but he is aware of his time inconsistency. Hence, he correctly believes that his future self will not maximize the current self’s future payoff. Therefore, in period one, the sophisticated consumer would like to commit to a smaller second period consumption. In the model, the shopping costs and the sales restriction provide him ways to commit.

How to differentiate between the sophisticated and the naive consumer? Consider state 0C. The sophisticated consumer would buy a larger inventory on Saturday than the naive consumer, in order to discourage his future self going back to the store on Sunday. Then on Sunday, in contrast to the naive consumer, he would never return to the store.

The Sunday sales restriction reduces the consumption of both the naive and the sophisticated consumer. More specifically, with the Sunday sales restriction, they will consume the same amount. This is because, the consumers are now essentially identical: they have the same preferences, and even though their beliefs about their future selves differ, the beliefs don’t matter anymore, as they cannot return to the store on Sunday. In contrast to the time consistent consumer, the sales restriction may increase the welfare of both the naive and the sophisticated consumer. This is because the reduction in consumption increases the welfare of both of them.

4 Identification and Estimation

4.1 Identification

To estimate the model I make the following assumptions.

Assumption 1. *Consumption shocks for beer η_{int} are independently and identically distributed across consumers i , weeks n , and days t , conditional on observable time and consumer characteristics, according to log-normal distribution.*

Specifically, I assume that $\log \eta_{int}$ is normally distributed with mean and standard deviation:

$$E[\log \eta_{int}] = \mu_0 + \mu_1 \cdot \text{Sunday}_{nt} + \mu_2 \cdot \text{Summer}_n + \mu_3 \cdot \text{Holidays}_n + \mu_4 \cdot \text{Inventory}_{in} \\ + \mu_5 \cdot \text{LogIncome}_{in} + \mu_6 \cdot \text{Single}_{in} + \mu_7 \cdot \text{NoHighSchool}_{in} + \mu_8 \cdot \text{GradSchool}_{in} \quad (6)$$

$$StDev [\log \eta_{int}] = \sigma_0 + \sigma_1 \cdot Sunday_{nt} \quad (7)$$

where $Sunday_{nt}$ is an indicator variable, $Summer_n$ is an indicator variable for dates from Memorial Day to Labor Day weekends, $Holidays_n$ is an indicator variable for holiday weekends, $Single_{in}$ is an indicator variable for single person households, $NoHighSchool_{in}$ and $GradSchool_{in}$ are both indicator variables describing household's highest education level, and $Inventory_{in}$ is an indicator variable for a beer purchase earlier during the same week, which is meant to capture possible beer inventory at home.

Regarding shopping costs I assume that it is observable whether shopping cost is zero or not.

Assumption 2. *If the consumer purchases a large quantity of other goods together with beer, then his shopping cost equals zero. Otherwise, shopping cost equals C .*

Note that C is a parameter that I will estimate. In the baseline model estimation, I assume that shopping cost equals zero if on a given day a consumer purchased goods other than beer worth at least as much as that paid for beer, and equals C otherwise.¹⁸ To check how sensitive the results are regarding the choice of the threshold, Appendix D presents estimates using two alternative thresholds: one third and two thirds.

Suppose that for each consumer for many pairs of time periods, we observe his quantity of beer purchased, prices (even when he did not purchase anything), and indicators for shopping costs constructed according to Assumption 2. Using the model and the data, we want to identify discount factor β , marginal utility from income α , health cost h , shopping cost C , and parameters of the consumption shock distribution. We can identify all these for each consumer separately, and I will describe below how to do that. However, when estimating the model I will assume that most parameters are the same across consumers.

The parameters βh and $\beta \alpha$ and of the distribution of the second period consumption shock F_2 are identified from the second period conditional demand in state C0 using data on different quantities and prices. In state C0, it is costly to purchase beer in the first period and costless in the second period. When the consumer didn't purchase anything in the first period then the second period demand does not depend on the first period consumption shock, because by assumption shocks are independent across periods. In particular, the conditional probability (conditioned on not purchasing in period one) of purchasing different quantities

¹⁸Note that the assumption of observable shopping costs can be at least partially relaxed, by assuming that if the consumer purchases a large quantity of other goods together with beer, then his shopping cost equals zero, and otherwise shopping cost equals zero with some probability.

at different prices identifies parameters βh , $\beta\alpha$, and F_2 :

$$Prob(x_2 > K|C0, x_1 = 0) = F_2 \left(\frac{1}{\beta(h + \alpha p)} - K \right), \quad K \geq 0 \quad (8)$$

Note that the identification does not rely on making any assumptions about the tie-breaking rule, i.e. consumer behavior in the situations when he is indifferent. That is the benefit of conditioning on $x_1 = 0$ in Equation 8. Specifically, conditioning on $x_1 = 0$ excludes the situations where the model predicts that the consumer is indifferent between purchasing inventory (while already at the store) and returning to the store in the next period. For the same reason, in the next steps it is useful to use the joint probability of no purchase in either of the periods. Then we don't need any assumptions about the tie-breaking rule in state 00.

The parameters of F_1 are identified from purchases in state 00, given βh , $\beta\alpha$, and parameters of F_2 that we know from above:

$$Prob(x_1 = 0, x_2 = 0|00) = \left[1 - F_1 \left(\frac{1}{\beta(h + \alpha p)} \right) \right] \left[1 - F_2 \left(\frac{1}{\beta(h + \alpha p)} \right) \right] \quad (9)$$

To identify discount factor β we compare purchases in states 00 and 0C.¹⁹ Specifically, by combining Equation (9) and

$$Prob(x_1 = 0|0C) = \left[1 - F_1 \left(\frac{1}{\beta(h + \alpha p)} \right) \right] \left[1 - F_2 \left(\frac{1}{h + \alpha p} \right) \right] \quad (10)$$

we can identify h , α , and β separately. The idea is that for a time consistent consumer the probability of not buying anything in these two states would be the same. The more time inconsistent a consumer is, the less likely he is to buy in advance in state 0C.

Given the rest of the parameters, shopping cost C is identified from the probability of purchasing in the first period in state C0:

$$Prob(x_1 > 0|C0) = F_1 \left(\frac{\Lambda(C)}{\beta(h + \alpha p)} \right) \quad (11)$$

because Λ is strictly monotone in C .

Note again that the identification in Equation 11 does not rely on the consumer not buying inventory in a situation where the model predicts that he is indifferent between doing so and returning to the store next period. When the consumer purchases something in the

¹⁹Alternatively, we could identify the discount factor from the probability of purchasing in the second period in state 0C.

first period, then the model predicts that he is indifferent between buying inventory now and returning to the store later. The above makes no predictions for that situation. However, planning to purchase inventory for period two cannot affect the probability of going shopping in period one, because shopping in period one is costly while it's costless in period two. Therefore, if the consumer doesn't buy anything in the first period, it must be only because of the value of his first period consumption shock.

Discussion of identification: Note that a crucial assumption in the identification is that there is no uncertainty regarding consumption shocks and shopping costs. That is, the consumer knows on Saturday his Sunday's shopping cost and consumption shock. Although it is a restrictive assumption, it is quite reasonable in this short time frame. However, if some consumers return to the store because of uncertainty and not because of time inconsistent discounting, then the model overestimates the share of time inconsistent consumers. What does that mean for the impact of the sales restriction? If there is uncertainty, the direct effect is that the sales restriction becomes more costly: it is more likely that the consumer has not stored enough, and due to the sales restriction he cannot return to the store. Because of this, the sales restriction might become less desirable.

It is also important to note that the model is estimated using household level data. I assume implicitly that a household behaves as a single decision maker. The assumption has non-trivial interaction with time-inconsistency. In particular, each individual in a household might behave as if time consistent, but when they have different objectives, then this might look like the household behaves as if time inconsistent. It is beyond the scope of this paper to model intra-household bargaining, and therefore I assume that the household members have the same objectives and information. A concern with this simplification is that I'm classifying some households as time inconsistent, where each individual separately behaves as if time consistent. This would affect the welfare implications.

The general idea how the model matches the data can be seen in Table 5, which presents the probability of purchasing beer in these four states. The comparison of the purchases in the state where the consumer buys only for current consumption (state 00) to that where he buys inventory (state 0C) should inform us about discounting. Similarly, shopping cost parameter C is captured by the difference of the probability of purchasing in period one in states where the shopping cost is zero and that where shopping is costly, namely states 00 and C0.

4.2 Estimation Strategy

I estimate the parameters of the model using the following data. For consumer i and over time period pairs indexed by $n = 1, \dots, N$, we observe beer prices $(p_{11}, p_{12}, \dots, p_{N1}, p_{N2})$, purchased beer quantities $(x_{11}, x_{12}, \dots, x_{N1}, x_{N2})$, and indicators for shopping costs constructed according to Assumption 2 $(C_{11}, C_{12}, \dots, C_{N1}, C_{N2})$.²⁰ Let's denote the above data by Y .

I estimate the share of time inconsistent consumers $1 - \rho$, their discount factor β , and other parameters $\theta = (\alpha, C, h, \mu, \sigma)$ from a mixture model by Maximum Likelihood. I formulate a mixture model by assuming each consumer is one of the two unobservable types, either time consistent or time inconsistent. That is, the consumer is time consistent (TC) with probability ρ and time inconsistent with probability $1 - \rho$. The likelihood equals:

$$L = \prod_i (\rho L_i(1, \theta, Y))^{1[i \text{ is } TC]} \cdot ((1 - \rho) L_i(\beta, \theta, Y))^{[1 - 1[i \text{ is } TC]]} \quad (12)$$

where $1[i \text{ is } TC]$ is an indicator function for time consistent consumer, $L_i(1, \theta, Y)$ and $L_i(\beta, \theta, Y)$ are consumer i 's contributions to the likelihood if he were time consistent and time inconsistent, respectively. The contribution of consumer i with a given discount factor to the likelihood is presented in Appendix C. After taking logarithms, I calculate an expectation of the log-likelihood, where the expectation is taken over the unobserved type. We can write the expectation of log-likelihood in the following way, where the unobserved type is replaced by the posterior probability that the consumer is time consistent:

$$E(\log L) = \sum_i [\rho_i \log \rho L_i(1, \theta, Y) + (1 - \rho_i) \log(1 - \rho) L_i(\beta, \theta, Y)] \quad (13)$$

where consumer i 's probability of being time consistent is calculated by Bayes' rule:

$$\begin{aligned} \rho_i &= E(1[i \text{ is } TC] | \rho, \beta, \theta, Y) = Prob_i(TC | \rho, \beta, \theta, Y) \\ &= \frac{\rho L_i(1, \theta, Y)}{\rho L_i(1, \theta, Y) + (1 - \rho) L_i(\beta, \theta, Y)} \end{aligned} \quad (14)$$

Thus, consumer i is time consistent with probability ρ_i and the likelihood of the time consistent consumer is ρ times his consumer-specific likelihood $L_i(1, \theta, Y)$, and similarly for the time inconsistent consumer. Note that the consumer-specific likelihood is the probability of observing the whole sequence of his purchases, not only the purchases from one time period.

²⁰According to the assumptions, it is observable when the consumer has shopping cost equal to zero. That is, he is in the store buying other goods, and can costlessly buy beer.

I estimate the model using the Expectation–Maximization (EM) algorithm. In the following, I will describe the EM algorithm estimation procedure. I maximize (13) over β , θ , and ρ using the following steps. First, I fix some initial values for β , θ , and ρ . Then, in step 2, given these parameter values and data for each consumer i , I calculate his probability of being time consistent, ρ_i , according to Equation (14) using Bayesian updating. In step 3, I substitute these consumer-specific probabilities ρ_i -s into log-likelihood Equation (13) to obtain new estimates of β , θ , and ρ . The maximization problem in step 3 can be separated in the following way:

$$(\hat{\beta}, \hat{\theta}) = \arg \max_{\beta, \theta} \sum_i [\rho_i \log L_i(1, \theta, Y) + (1 - \rho_i) \log L_i(\beta, \theta, Y)] \quad (15)$$

$$\hat{\rho} = \frac{1}{I} \sum_{i=1}^I \rho_i \quad (16)$$

where I denotes the number of consumers. Using the new parameter estimates from (15) and (16), I go back to step 2 and calculate new ρ_i -s. I iterate over steps 2 and 3 until convergence.

5 Estimation Results

Table 6 presents the parameter estimates. The baseline model in Column 1 gives an estimate of the share of time consistent consumers ρ equal to 0.84 and the time inconsistent discount factor β equal to 0.47. Both of these estimates are significantly below one.

My estimate of the time inconsistent discount factor is in the range that previous studies have found (for an overview, see Frederick, Loewenstein, and O’Donoghue (2002)). Most studies have estimated a common discount factor for all consumers, while I allow two types of consumers, either with discount factor equal to 1 or with discount factor $\beta < 1$. Hence, my estimate of β is the discount factor for the especially present biased group of consumers. Recent studies using data from various topics estimate a common time inconsistent discount factor to be from 0.5 to 0.9. For example, based on savings and consumption decisions, Laibson, Repetto, and Tobacman (2007) estimate $\beta = 0.7$, and Brown, Chua, and Camerer (2009) from 0.6 to 0.7; Fang and Wang (2015) estimate β to be from 0.56 to 0.71 based on mammography decisions; Tanaka, Camerer, and Nguyen (2010) in experiments in Vietnam get β to be from 0.74 to 0.89.

Although only a small percentage of consumers are estimated to be time inconsistent, they account for a larger share in total beer consumption. According to the estimates from the

baseline model, time inconsistent consumers form 16% of the sample of frequent beer buyers and only 3% of the full sample. The model predicts that these 16% of consumers account for 52% of beer consumption in the sample of frequent beer buyers. In the full sample, these 3% of consumers account for 40% of total beer consumption. This drives the welfare implications.

For time inconsistent consumers, discount factor β directly corresponds to the imaginary price reduction. That is, for time inconsistent consumers the effective price is only about half of the actual price (when he buys for the current period). This also implies that if the government could discriminate and tax time inconsistent consumers, it could offset the present biased discounting.

Price elasticity of beer according to my estimates equals - 2.2 for time consistent consumers and -1.8 for time inconsistent consumers (see Table 7).²¹ The estimated elasticities are somewhat higher than in the literature, where alcohol demand is often found inelastic. The estimates imply that the time inconsistent consumers respond less to price. Recalling the discussion in the previous paragraph, this is exactly what we would expect.

Estimated shopping cost from the baseline model is slightly above five dollars. Remember that the shopping cost can take two values, either zero or this estimated higher value, and therefore the average shopping cost is lower. There are somewhat similar shopping cost values estimated in the literature. For example, Erdem, Imai, and Keane (2003) estimate ketchup demand and find that for the most popular size (32 ounce) ketchup, their estimated shopping cost is slightly higher than the typical price of the product.

In Columns 2-4 in Table 6, several variables are added to the beer consumption shock parameter μ , according to Equation (6). Column 2 adds household characteristics. Column 3 includes indicator variables for holidays, summer, beer purchases during the weekdays of the same week (a proxy for inventory), and health cost parameter h , but only for consumers with a college degree, modeled as $h \cdot IndicatorCollege$. Column 4 combines all the variables. The main parameter estimates are similar across all the models.

To check the robustness of the estimation results, the model in Column 1 in Table 6 is re-estimated with modified datasets; Table 9 in Appendix D presents the results. In Column 1 of Table 9, holidays are excluded from the sample. In Column 2, the sample excludes weeks with previous beer purchases during the weekdays. In Column 3, the sample includes only single person households. In Column 4, the sample includes only households who always purchase beer at either grocery stores or discount stores. In Column 5, the sample includes only the households who purchase only smaller quantities of beer (up to 18

²¹The elasticity is calculated numerically using the predicted demand function at the median price.

12-ounce bottles/cans per day). Datasets in Columns 6 and 7 use alternative thresholds for shopping cost. In the main dataset, shopping cost is assumed to equal zero if, on a given day, a consumer purchased goods other than beer worth at least as much as that paid for beer, and to equal C otherwise. In Columns 6 and 7, the thresholds for the share of beer expenditures from total expenditures equal one third and two thirds, respectively. In these estimates in Table 9, the share of time consistent consumers is in the range of 0.74 to 0.90 and discount factor from 0.33 to 0.61.

Finally, I estimate a restricted version of the model with the households that were excluded from the main sample. The main sample was restricted to the households that often purchase beer. However, to understand the implications of the regulation, we should know how it affects the whole population. I assume that the infrequent consumers are otherwise similar to the consumers in the main sample, except that they are all time consistent and don't want to consume much beer. Using a dataset of the infrequent consumers, I estimate a restricted model, where I restrict all the parameter values to equal to the parameter estimates from Column 1 in Table 6 except the parameters of the consumption shocks distributions. The results are in Column 5 in Table 6.

6 Counterfactual Analysis

In this section, I discuss the policy analysis. Specifically, I look at what happens to consumer welfare when the Sunday sales restriction is introduced or taxes are increased. In the sample of frequent consumers, the parameter estimates used in the counterfactual analysis are from Column 1 in Table 6. To extrapolate the results to the full sample, i.e. to include the infrequent consumers, I use the parameter estimates in Column 5 in Table 6 obtained from the restricted model with the data of infrequent consumers. Whenever applicable, I extrapolate the weekend results to the whole week under the assumption that the sales restriction does not affect purchases on weekdays and the share of weekend purchases equals 0.367.

6.1 Sunday Sales Restriction

Before analyzing changes in welfare, let us first consider how the sales restriction impacts consumption. Figure 6a presents the time consistent consumer's percentage decrease in total consumption as a function of shopping cost parameter C with and without the sales restriction. Without shopping costs, the time consistent consumer would buy in advance and store the same quantity that he otherwise would have bought in the second period. With

shopping costs, his total demand with the sales restriction is lower, due to those weeks when it is costly for him to shop on Saturdays.

Even without any shopping costs, the sales restriction decreases the time inconsistent consumer's consumption because when he has to buy in advance he buys less. To see how the decrease in consumption is related to the first period shopping cost, see Figure 6b. Note that the impact of the restriction is larger with small shopping costs compared to large shopping costs. This is because a large shopping cost itself provides a partial commitment device. There is a threshold below which shopping costs are too low to prevent the consumer from returning to the store. Therefore, below the threshold, the impact of the sales restriction is larger.

How much the sales restriction decreases consumption (at the estimated shopping cost) is presented in Column 3 in Table 8. Time consistent consumers will mostly buy in advance, so that their consumption decreases only by 3%. As expected, the impact is larger on time inconsistent consumers; their consumption decreases by 20%.

The aggregate decrease in consumption depends on the impact on each consumer group as well as the relative sizes of the groups. Recalling that the main sample is restricted to frequent consumers (who buy beer often), I will present back-of-the-envelope calculations which extrapolate the results to the full sample. The model predicts that the sales restriction decreases infrequent consumers' consumption by 5%. Combining the results from the frequent and infrequent consumers, the total consumption decreases by 10%. This is in line with but slightly larger than the 8% increase in beer sales in Sweden after the policy experiment of lifting sales bans (Norström and Skog, 2003, 2005).

For the welfare analysis, let me first clarify how I define the welfare of the time inconsistent consumer. In most of the analysis, following a common approach in the literature (see for example O'Donoghue and Rabin (2006)), I use the welfare of the long-term self, which is the time inconsistent consumer's utility function without time inconsistent discounting. That is, time inconsistent discount factor β determines the behavior of the consumer and through that indirectly affects his welfare, but does not directly enter the welfare function. This is a paternalistic view that treats the time inconsistent consumer's preferences for immediate consumption as an error. With this welfare criterion, the social welfare function of a heterogeneous population weighs equally all consumers. In addition to the above, I will also consider two alternative welfare criteria. First, what if the social planner does not weigh all the consumers equally? Specifically, I ask what should the welfare weight of time inconsistent consumers be in the social welfare function for the sales restriction to be optimal?

In the second alternative approach, I calculate the social welfare, where the time inconsistent consumer's welfare is his utility in period $t = 0$. That is, the welfare function weighs all consumers equally, but the time inconsistent consumer discounts by β all his future payoffs. The welfare function might capture the voters' preferences.

Starting with the welfare criterion of the long term self of the time inconsistent consumer, let us look how the sales restriction affects consumer welfare. Recall that the restriction always decreases time consistent consumer welfare, as long as shopping costs are not zero. Thus, for the sales restriction to improve aggregate consumer welfare, the share of time inconsistent consumers and the commitment effect for time inconsistent consumers has to be large enough to compensate for the welfare loss from the increased inconvenience, both to the time inconsistent consumers themselves and also overall to the time consistent consumers. Figure 7a illustrates that in the sample of frequent consumers, the sales restriction improves aggregate consumer welfare as long as the share of time inconsistent consumers is larger than 5.8%. The estimated share of time inconsistent consumers in the sample of frequent consumers is 16%, which is above the threshold that makes the sales restriction welfare improving.

Since the main sample was restricted to 20% of the consumers, I will again extrapolate the welfare impact of the sales restriction on the full sample of households. By assumption the infrequent consumers are time consistent, therefore the restriction decreases their welfare. However, they consume little, so we would expect the impact to the total welfare to be relatively small. Recall from Table 8 that the share of time inconsistent consumers in the full sample is 3%. According to Figure 7b, the sales restriction improves consumer welfare in the full sample as long as the share of time inconsistent consumers is larger than 1.8%. Hence, the sales restriction is still welfare improving. Note that the result is subject to the caveats that it is based on only approximations and relies on strong assumptions about the behavior of the households that are excluded from most of the analysis.

Alternatively, we could ask what welfare weight should the social planner place on time inconsistent consumers for the sales restriction to be optimal? Figure 7c presents the change in welfare due to the sales restriction at different weights on time inconsistent consumers while keeping the weight on time consistent consumers constant and equal to one. The figure presents the welfare change that is extrapolated to the full sample. As long as the social planner has a weight larger than 0.56 on each time inconsistent consumer, the sales restriction improves consumer welfare. Interestingly, the time inconsistent consumer values his own future welfare even less as he discounts it using the estimated discount factor equal to 0.47. I use this

welfare measure to take another look at the impact of the sales restriction. Figure 7d presents the change in welfare due to the sales restriction, where the time inconsistent consumer's welfare is measured by his utility in period $t = 0$ and the calculations are extrapolated to the full sample. Under this welfare measure the time inconsistent consumer values his future payoff much less than the time consistent consumer. In this way it is equivalent to the social planner placing a much lower weight on the time inconsistent consumer. Figure 7d shows that under this welfare measure the sales restriction becomes welfare decreasing because only 3% of the consumers are time inconsistent, which is below the positive impact threshold of the restriction. Recall that the sales restriction can benefit only time inconsistent consumers, it always hurts time consistent consumers. As time inconsistent consumers place such a low weight on their own future payoffs (or equivalently, the social planner places a low weight on time inconsistent consumers' welfare), there are not enough time inconsistent consumers for the sales restriction to be beneficial.

6.2 Comparison with Sales Tax

Alcohol is regulated mainly because it is considered to be a good with negative externality. The externality includes both direct costs for the society like drunk driving and indirect costs like alcohol consumers' health problems. Due to these costs the socially optimal level of alcohol consumption is lower than the private optimum. The current paper does not consider the externality from alcohol and cannot find the optimal alcohol consumption level taking into account the externality. But, we can ask the question of which policy instrument we should use to decrease consumption.

In the standard model, economic theory tells us that we can obtain the first best solution using Pigouvian taxes, distortionary taxes that equate the social payoff with the private one. When consumers are time inconsistent, are taxes still welfare improving compared to the sales restriction? We saw in the previous section that the sales restriction decreases the welfare of time consistent consumers and increases the welfare of time inconsistent consumers. We can expect taxes to have similar effects: time consistent consumers won't like taxes, while taxes reduce time inconsistent consumers' consumption which may increase welfare. However, the sales restriction might be preferred to taxes, because the sales restriction is a way to target the regulation. It doesn't have much effect on time consistent consumers, while it has a large impact reducing the consumption of time inconsistent consumers.

In the following, I will look at whether taxes are welfare improving compared to the sales restriction in achieving an equivalent decrease in consumption. First, I look at the aggregate

consumption and find the tax rate that leads to the equivalent consumption decrease as the sales restriction. These tax rates, for each consumer group separately and for the full sample, are presented in Table 8. Then I compare the welfare impact of this tax rate and the sales restriction. Note that the welfare from taxes is calculated assuming that the tax revenue is redistributed to the consumers as a lump-sum payment.

Let's first consider the welfare of the time consistent consumer. Figure 8a describes the decrease in the time consistent consumer's welfare due to these policies. Both the sales restriction and taxes decrease the time consistent consumer's welfare. Suppose all consumers were time consistent and suppose we want to reduce their alcohol consumption by a given amount. Should we then restrict sales or increase taxes? Taxes are more flexible in allowing time consistent consumers to choose when to decrease consumption; therefore, time consistent consumers would prefer taxes. Figure 8a illustrates that the time consistent consumer prefers taxes to the sales restriction at lower tax rates and sales restriction to taxes at higher tax rates. We are interested in the comparison of the policy measures which give the same decrease in consumption. According to Table 8 the sales restriction leads to a decrease in consumption that is equivalent to a tax rate increase of 1.2 percentage points. Figure 8a demonstrates that this tax increase is welfare improving compared to the sales restriction.

Figure 8b describes the welfare of the time inconsistent consumer showing that both policies improve his welfare. According to Table 8 the sales restriction decreases the time inconsistent consumer's consumption by 20% and the same reduction in consumption can be achieved by a tax rate increase of 9.5 percentage points. Figure 8b illustrates that, in terms of consumer welfare, this tax increase is preferable to the sales restriction. As discussed above, the sales restriction has two effects. First, for both types of consumers, it increases costs by making it more inconvenient to buy. But, for time inconsistent consumers, it also has a positive effect, as it provides them a commitment device. In comparison, taxes are more flexible in allowing consumers to choose when to decrease consumption. Since time consistent consumers do not benefit from commitment, they always prefer the tax increase. According to the estimates, the commitment effect for the time inconsistent consumers is not large enough to compensate for the welfare loss from the sales restriction, therefore they also prefer the tax increase.

Figures 8c and 8d present the aggregate consumer welfare with the sales restriction and taxes in the sample of frequent consumers and full sample respectively. Here a common tax rate is applied to all consumers and the tax revenue distributed back to the consumers as a lump sum payment. Recall from the above that for time consistent consumer to prefer taxes

to the sales restriction, the tax increase must be low enough, whereas for time inconsistent consumer to prefer taxes, the tax increase must be high enough. Interestingly, the common tax increase that is equivalent to the sales restriction is so small that in terms of the time inconsistent consumer's welfare the sales restriction would be preferable to the tax increase. However, in terms of the aggregate consumer welfare the tax increase is preferable to the sales restriction.

7 Conclusions

I provided a model where a sales restriction decreases consumption if consumers are time inconsistent or shopping is costly. Using scanner data of beer purchases, I estimated the fraction of time inconsistent consumers to be 16% among the regular beer buyers or 3% among all consumers. I found that the sales restriction may increase aggregate consumer welfare. It provides a commitment device for time inconsistent consumers that allows them to reduce their consumption. However, when the goal is to reduce alcohol consumption, then taxes would be welfare improving compared to the sales restriction.

The result supports the recent trend of relaxing sales restrictions. The trend is likely to reflect that either reducing alcohol consumption is not optimal or there are better policy tools for that than the sales restriction. Whether or not reducing alcohol consumption is socially beneficial is out of the scope of the paper because I don't explicitly consider the reasons for reducing alcohol consumption, nor the externalities associated with alcohol consumption. Finally, it is important to note that when evaluating the optimality of the sales restriction, I restricted attention to only one possible cost – the inconvenience of costly shopping trips. However, the other potential costs related to the restriction might outweigh the benefits.

References

- BANERJEE, A., AND S. MULLAINATHAN (2010): "The Shape of Temptation: Implications for the Economic Lives of the Poor," *NBER Working Paper*, No. 15973.
- BERNHEIM, B. D., J. MEER, AND N. K. NOVARRO (2016): "Do Consumers Exploit Precommitment Opportunities? Evidence from Natural Experiments Involving Liquor Consumption," *American Economic Journal: Economic Policy*, p. forthcoming.

- BERNHEIM, B. D., AND A. RANGEL (2004): “Addiction and Cue-Triggered Decision Processes,” *The American Economic Review*, 94(5), 1558–1590.
- BESHEARS, J., J. J. CHOI, D. LAIBSON, AND B. MADRIAN (2006): “Early Decisions: A Regulatory Framework,” *NBER Working Paper*, No. 11920.
- BROWN, A. L., Z. E. CHUA, AND C. F. CAMERER (2009): “Learning and Visceral Temptation in Dynamic Saving Experiments,” *The Quarterly Journal of Economics*, 124(1), 197–231.
- CARPENTER, C. S., AND D. EISENBERG (2009): “Effects of Sunday Sales Restrictions on Overall and Day-Specific Alcohol Consumption: Evidence From Canada,” *Journal of Studies on Alcohol and Drugs*, 70(1), 126–133.
- DELLAVIGNA, S., AND U. MALMENDIER (2006): “Paying Not to Go to the Gym,” *American Economic Review*, 96(3), 694–719.
- ERDEM, T., S. IMAI, AND M. P. KEANE (2003): “Brand and Quantity Choice Dynamics Under Price Uncertainty,” *Quantitative Marketing and Economics*, 1(1), 5–64.
- FANG, H., AND D. SILVERMAN (2009): “Time-inconsistency and Welfare Program Participation: Evidence from the NLSY,” *International Economic Review*, 50(4), 1043–1077.
- FANG, H., AND Y. WANG (2015): “Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions,” *International Economic Review*, 56(2), 565–596.
- FREDERICK, S., G. LOEWENSTEIN, AND T. O’DONOGHUE (2002): “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 40(2), 351–401.
- GRUBER, J., AND D. M. HUNGERMAN (2008): “The Church Versus the Mall: What Happens When Religion Faces Increased Secular Competition?,” *The Quarterly Journal of Economics*, 123(2), 831–862.
- GRUBER, J., AND B. KÖSZEGI (2001): “Is Addiction “Rational”? Theory and Evidence,” *The Quarterly Journal of Economics*, 116(4), 1261–1303.
- (2004): “Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes,” *Journal of Public Economics*, 88(9-10), 1959–1987.

- HAAVIO, M., AND K. KOTAKORPI (2011): “The political economy of sin taxes,” *European Economic Review*, 55(4), 575–594.
- HASTINGS, J., AND E. WASHINGTON (2010): “The First of the Month Effect: Consumer Behavior and Store Responses,” *American Economic Journal: Economic Policy*, 2(2), 142–162.
- HENDEL, I., AND A. NEVO (2006): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74(6), 1637–1673.
- (2013): “Intertemporal Price Discrimination in Storable Goods Markets,” *The American Economic Review*, 103(7), 2722–2751.
- LAIBSON, D. (1997): “Golden Eggs and Hyperbolic Discounting,” *The Quarterly Journal of Economics*, 112(2), 443–478.
- LAIBSON, D., A. REPETTO, AND J. TOBACMAN (2007): “Estimating Discount Functions with Consumption Choices over the Lifecycle,” *NBER Working Paper*, No. 13314.
- MARCUS, J., AND T. SIEDLER (2015): “Reducing binge drinking? The effect of a ban on late-night off-premise alcohol sales on alcohol-related hospital stays in Germany,” *Journal of Public Economics*, 123, 55–77.
- MASTROBUONI, G., AND M. WEINBERG (2009): “Heterogeneity in Intra-Monthly Consumption Patterns, Self-Control, and Savings at Retirement,” *American Economic Journal: Economic Policy*, 1(2), 163–189.
- MEIER, S., AND C. SPRENGER (2010): “Present-Biased Preferences and Credit Card Borrowing,” *American Economic Journal: Applied Economics*, 2(1), 193–210.
- NORSTRÖM, T., AND O.-J. SKOG (2003): “Saturday opening of alcohol retail shops in Sweden: an impact analysis,” *Journal of studies on alcohol*, 64(3), 393–401.
- (2005): “Saturday opening of alcohol retail shops in Sweden: an experiment in two phases,” *Addiction*, 100(6), 767–776.
- O’DONOGHUE, T., AND M. RABIN (2006): “Optimal sin taxes,” *Journal of Public Economics*, 90(10-11), 1825–1849.

- PASERMAN, M. D. (2008): “Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation*,” *The Economic Journal*, 118(531), 1418–1452.
- POPOVA, S., N. GIESBRECHT, D. BEKMURADOV, AND J. PATRA (2009): “Hours and Days of Sale and Density of Alcohol Outlets: Impacts on Alcohol Consumption and Damage: A Systematic Review,” *Alcohol and Alcoholism*, 44(5), 500–516.
- SHAPIRO, J. M. (2005): “Is there a daily discount rate? Evidence from the food stamp nutrition cycle,” *Journal of Public Economics*, 89(2-3), 303–325.
- STEHR, M. (2007): “The Effect of Sunday Sales Bans and Excise Taxes on Drinking and Cross-Border Shopping for Alcoholic Beverages,” *National Tax Journal*, 60(1), 85.
- TANAKA, T., C. F. CAMERER, AND Q. NGUYEN (2010): “Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam,” *American Economic Review*, 100(1), 557–571.
- TAROZZI, A., AND A. MAHAJAN (2011): “Time Inconsistency, Expectations and Technology Adoption: The Case of Insecticide Treated Nets,” *Unpublished manuscript, Stanford University*.

Tables and Figures

Table 1: Summary statistics of households

	HH in Sample	Other HH
Number of households	5385	22034
Median age	50 - 54	50 - 54
Median income (\$)	50,000 - 59,999	50,000 - 59,999
Median family size	2	2
Share of HH incl. non-fulltime working male	0.184	0.148
Share of HH incl. fulltime working female	0.381	0.381
Share of HH with college education	0.457	0.531
Share of HH with children	0.260	0.300
Median number of days with beer purchases per year	18	2
Median expenditure on beer per year (\$)	244.42	18.45

Note: An observation is a household.

Table 2: Shopping trip characteristics of the households in the main sample

	Mean	Sd	Min	Max	Median
Total expenditure (\$)	77.920	93.499	0.000	3297.100	49.570
Share of beer in expenditures	0.081	0.179	0.000	1.000	0.000
Quantity of beer	0.942	0.871	0.024	36.750	0.750
Price of beer(\$)	16.899	1.638	12.494	21.918	16.787
Total beer expenditure (\$)	15.613	13.471	0.240	673.680	12.880

Note: For total expenditure and share of beer, an observation is a shopping trip on weekends. For other shopping trip characteristics, an observation is a shopping trip with a beer purchase on weekends. The measurement unit of the quantity of beer is an equivalent to a case of 24 12-ounce cans/bottles. The price is calculated as a weighted average of the prices of all beer sales in the state in each week; the prices are weighted by the sales. The sample is restricted to the households who frequently buy beer.

Table 3: Frequency of beer purchases

	(1)	(2)	(3)
log(Income)	-0.392*** (0.076)	-0.177 (0.128)	-0.194 (0.152)
Female	-0.656*** (0.196)	-0.759*** (0.247)	
Education: less than high school	0.795** (0.337)	-0.164 (0.251)	-0.391 (0.288)
Education: college	0.095 (0.136)	-0.089 (0.251)	-0.104 (0.299)
Education: graduate school	-0.465*** (0.161)	-0.340 (0.301)	-0.299 (0.336)
Age: under 40	-0.320* (0.170)	-0.171 (0.367)	-0.197 (0.451)
Age: above 65	-0.626*** (0.102)	-0.409 (0.283)	-0.426 (0.351)
Female working full time	0.278*** (0.106)	-0.341* (0.204)	
Male not working full time	0.000 (0.119)	-0.091 (0.333)	-0.099 (0.371)
Children	0.080 (0.247)		
Household size: 1 person	0.110 (0.165)		
Household size: 3-4 persons	-0.028 (0.133)		
Household size: 5 or more persons	-0.024 (0.314)		
Week fixed effects	Yes	Yes	Yes
Day of week fixed effects	Yes	Yes	Yes
Number of observations	494076	95778	67125

Note: Each columns presents estimates from a separate regression. Dependent variable is an indicator for purchasing beer given a beer purchase the previous day. The sample includes only the frequent consumers. In column 2, the sample is further restricted to single person households and in column 3, to single men. Standard errors are in parantheses. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table 4: Quantity of beer purchases

	(1)	(2)	(3)
log(Income)	-0.363*** (0.077)	-0.081 (0.059)	-0.044 (0.083)
Female	-0.633*** (0.189)	-0.095 (0.118)	-0.250 (0.169)
Education: less than high school	0.864*** (0.332)	-0.643* (0.344)	-1.147** (0.452)
Education: college	0.044 (0.125)	0.312*** (0.094)	0.413*** (0.159)
Education: graduate school	-0.539*** (0.159)	0.381*** (0.112)	-0.203 (0.189)
Age: under 40	-0.357** (0.163)	0.153 (0.110)	-0.107 (0.219)
Age: above 65	-0.607*** (0.101)	-0.104 (0.090)	-0.081 (0.131)
Female working full time	0.237** (0.099)	0.128* (0.076)	0.153 (0.117)
Male not working full time	0.004 (0.116)	-0.005 (0.088)	0.003 (0.132)
Children	0.029 (0.217)	0.221 (0.145)	0.229 (0.316)
Household size: 1 person	0.026 (0.164)	0.418*** (0.109)	0.210 (0.147)
Household size: 3-4 persons	-0.028 (0.125)	0.084 (0.106)	-0.048 (0.186)
Household size: 5 or more persons	-0.017 (0.284)	-0.208 (0.184)	-0.275 (0.350)
log(Quantity previous day)	-0.439*** (0.060)		
Beer purchase previous day		0.400*** (0.111)	
Six-pack pervious day			2.302*** (0.112)
Week fixed effects	Yes	Yes	Yes
Day of week fixed effects	Yes	Yes	Yes
Number of observations	494076	494076	60841

Note: Each columns presents estimates from a separate regression. In column 1, the dependent variable is an indicator for purchasing beer given a beer purchase the previous day. In columns 2 and 3, the dependent variable is an indicator for purchasing a six-pack given a beer purchase. The sample includes only the frequent consumers. In column 3, the sample is further restricted to days when household purchased beer the previous day. Standard errors are in parantheses. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table 5: Probability of purchasing beer in four states

	00	C0	0C	CC
Period 1	0.255	0.026	0.298	0.025
Period 2	0.188	0.302	0.022	0.021
Periods 1 and 2	0.059	0.007	0.008	0.002
Neither 1 nor 2	0.616	0.679	0.688	0.957
Number of Obs	91205	190291	171896	387101

Note: Sample is restricted to the households who frequently buy beer. An observation is a household-weekend pair. Each column represents one of the four states defined by shopping cost pairs over the weekend.

Table 6: Structural model parameter estimates

	(1)	(2)	(3)	(4)	(5)
Share of time consistent consumers ρ	0.840 (0.009)	0.841 (0.010)	0.864 (0.014)	0.836 (0.016)	
Discount factor β	0.474 (0.029)	0.465 (0.032)	0.396 (0.020)	0.434 (0.028)	
High shopping cost C^H	0.220 (0.041)	0.231 (0.045)	0.323 (0.014)	0.269 (0.039)	
Marginal utility of income α	0.042 (0.002)	0.044 (0.002)	0.041 (0.002)	0.042 (0.002)	
Health cost h			0.070 (0.068)	0.059 (0.013)	
μ_0	1.437 (0.336)	1.620 (0.434)	1.765 (0.077)	1.868 (0.522)	7.944 (1.320)
σ_0	0.832 (0.195)	0.862 (0.219)	1.079 (0.046)	0.935 (0.260)	2.738 (0.433)
Sunday μ_1	-0.264 (0.347)	-0.332 (0.383)	-0.574 (0.069)	-0.387 (0.472)	-5.205 (1.316)
Sunday σ_1	0.321 (0.193)	0.363 (0.200)	0.182 (0.065)	0.279 (0.240)	-1.356 (0.429)
Inventory μ_4			0.041 (0.027)	0.051 (0.029)	
Log Income μ_5		-0.011 (0.021)		-0.027 (0.018)	
Single μ_6		-0.026 (0.041)		-0.032 (0.038)	
No High school μ_7		-0.304 (0.110)		-0.282 (0.116)	
Graduate School μ_8		0.131 (0.035)		0.104 (0.034)	
Seasonal effects	No	No	Yes	Yes	No

Note: In columns 1 - 4, the sample is restricted to households who frequently buy beer. In column 5, the sample is restricted to infrequent consumers. *Seasonal effects* include indicators for holidays and summer. In column 5, the rest of the parameter values are set equal to the parameter estimates from Column 1. Bootstrapped standard errors are included in parenthesis.

Table 7: Demand elasticities

	00	0C	C0	CC	Aggregate
Time consistent	-2.607	-2.607	-2.305	-1.556	-2.176
Time inconsistent	-2.223	-1.929	-1.885	-1.401	-1.778

Note: The elasticities are calculated numerically using the predicted demand function at the median price and using the parameter estimates from column 1 in Table 6.

Table 8: Counterfactuals: the impact of the sales restriction on consumption

	Sample share (%)	Quantity share (%)	Quantity decrease (%)	Equivalent tax increase
Sample of frequent consumers	19.6	76.5	12.1	5.1
.. Time consistent	16.5	36.9	3.2	1.2
.. Time inconsistent	3.1	39.6	20.4	9.5
Sample of infrequent consumers	80.4	23.5	4.7	1.7
Full sample	100.0	100.0	10.3	4.2

Note: Sample share is the share of the consumer group in the full sample. Quantity share is the predicted quantity share of the consumer group. Quantity decrease is the predicted decrease of quantity due to the sales restriction. Quantity decrease due to sales restriction is extrapolated to the whole week. Equivalent tax increase is the percentage point tax increase that leads to the quantity decrease given in the previous column. Counterfactual analysis uses parameter estimates in Column 1 and 5 in Table 6.

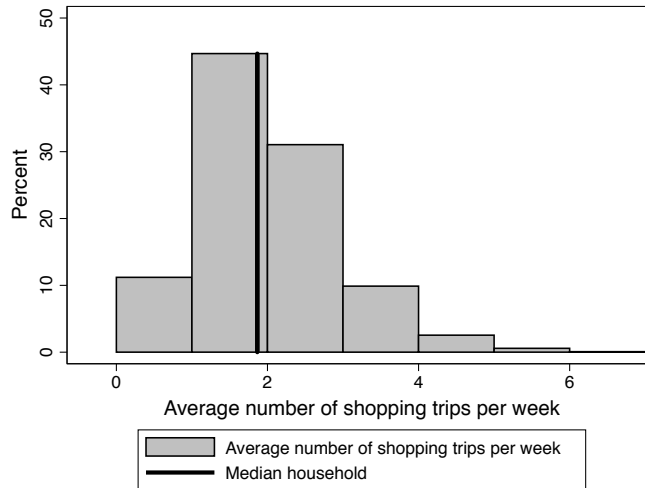


Figure 1: Number of shopping trips per week

Note: Data is from the main sample of consumers who buy beer frequently. An observation is a household-week pair. Shopping trips are aggregated to a day level.

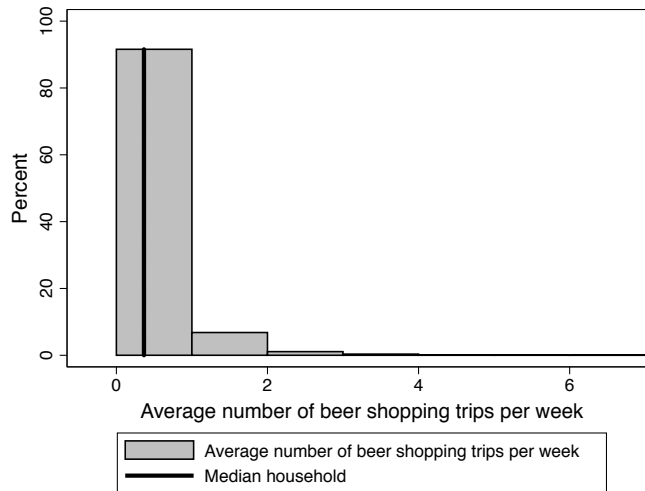


Figure 2: Number of beer shopping trips per week

Note: Data is from the main sample of consumers who buy beer frequently. An observation is a household-week pair. Shopping trips are aggregated to a day level.




Figure 3: Distribution of beer expenditures over the days of the week

Note: Data is from the main sample of consumers who buy beer frequently.

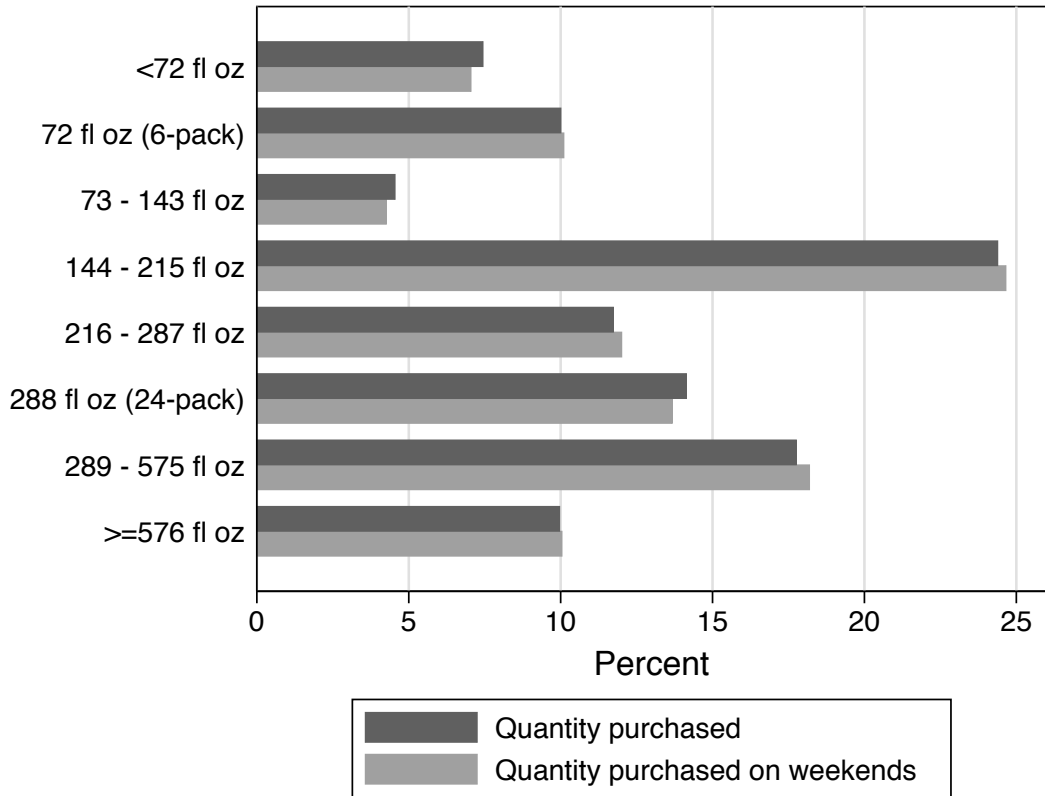


Figure 4: Histogram of beer quantity bought per shopping trip

Note: Data is from the main sample of consumers who buy beer frequently.

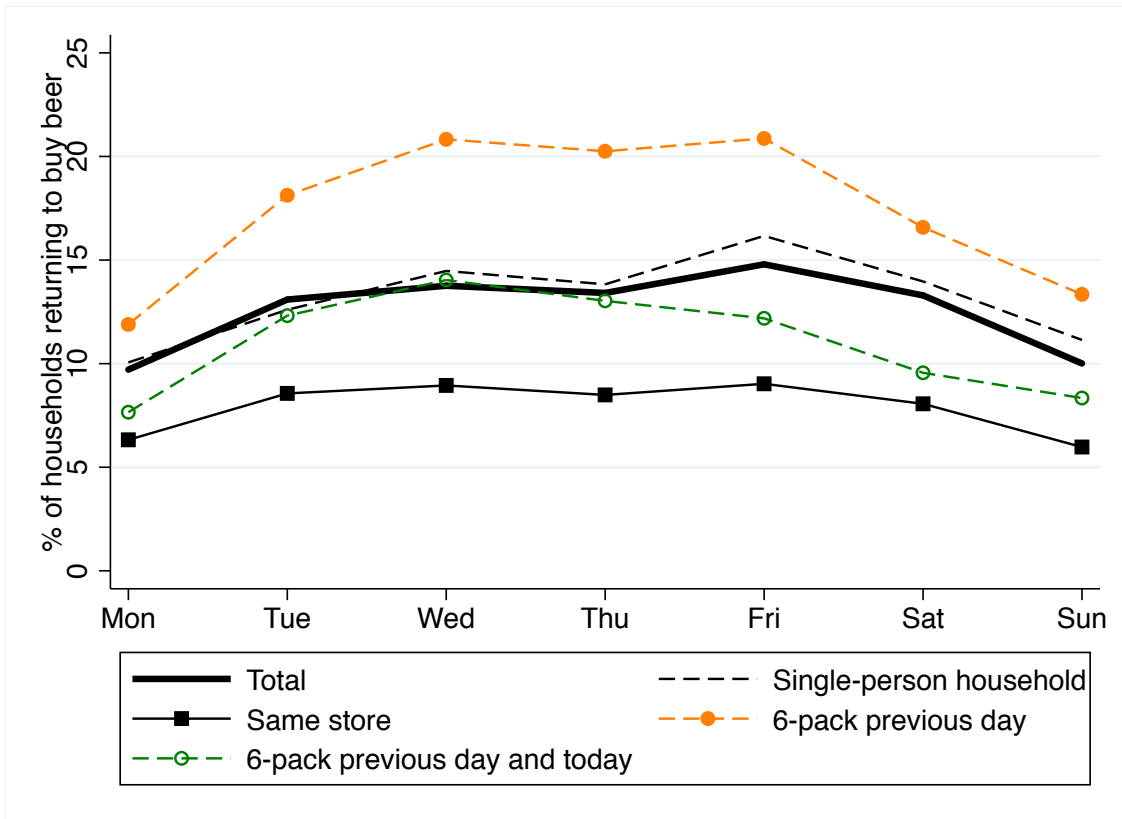

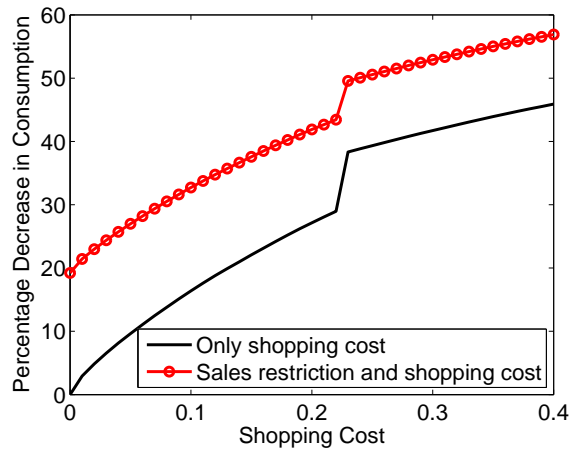


Figure 5: Percentage of households buying beer, given that they bought beer on the previous day, over the days of the week

Note: Data is from the main sample of consumers who buy beer frequently.



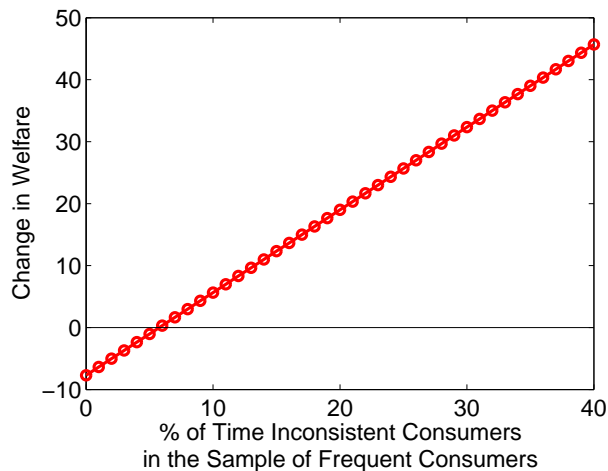
(a) Time consistent



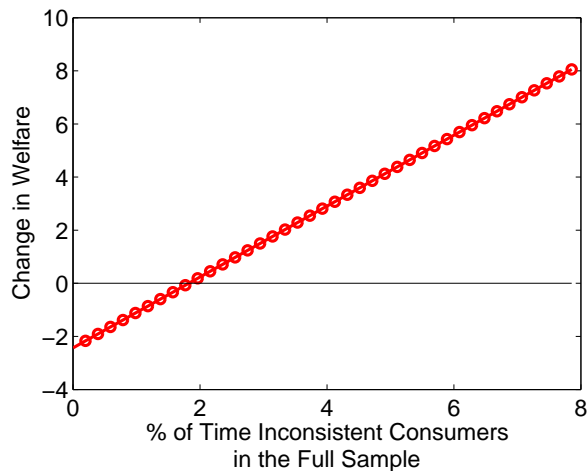
(b) Time inconsistent

Figure 6: Decrease in consumption due to shopping costs and the sales restriction

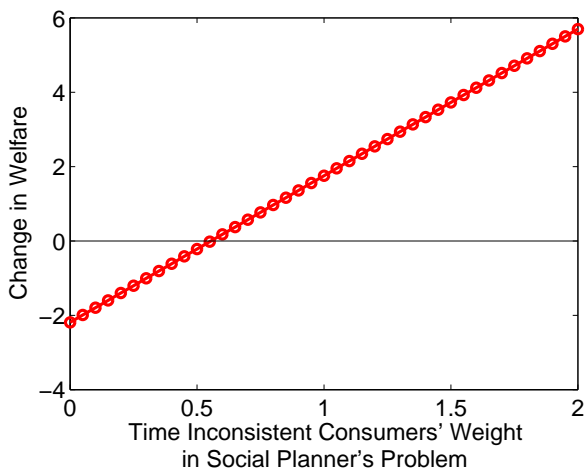
Note: Counterfactual analysis uses parameter estimates in Column 1 in Table 6 and data from the sample of frequent consumers. The percentage decrease in consumption is extrapolated (from the weekend consumption) by making the following assumptions: (a) shopping costs affect the weekend and weekdays equally, (b) sales restriction affects purchases only on the weekend. To extrapolate the decrease in consumption from the sales restriction, the percentage decrease in the weekend consumption is adjusted by the share of weekends in total sales, which equals 0.367.




(a) Sample of frequent consumers



(b) Full sample



(c) Time inconsistent consumers' weight in social planner's problem, full sample



(d) Time inconsistent consumers' lower value of future welfare, full sample

Figure 7: Change in aggregate consumer welfare due to the sales restriction

Note: Subfigure (a) uses parameter estimates from column 1 in Table 6 and data from the sample of frequent consumers. Subfigures (b)-(d) use parameter estimates from columns 5 in Table 6 in addition to the above and data from the full sample. On subfigures (a) - (c), time inconsistent consumer's welfare is his 'long term self's' welfare (utility function without time inconsistent discounting β). On subfigure (d), time inconsistent consumer's welfare is measured by his utility in period $t = 0$. On y-axes, the change in welfare is measured in dollars and annualized.

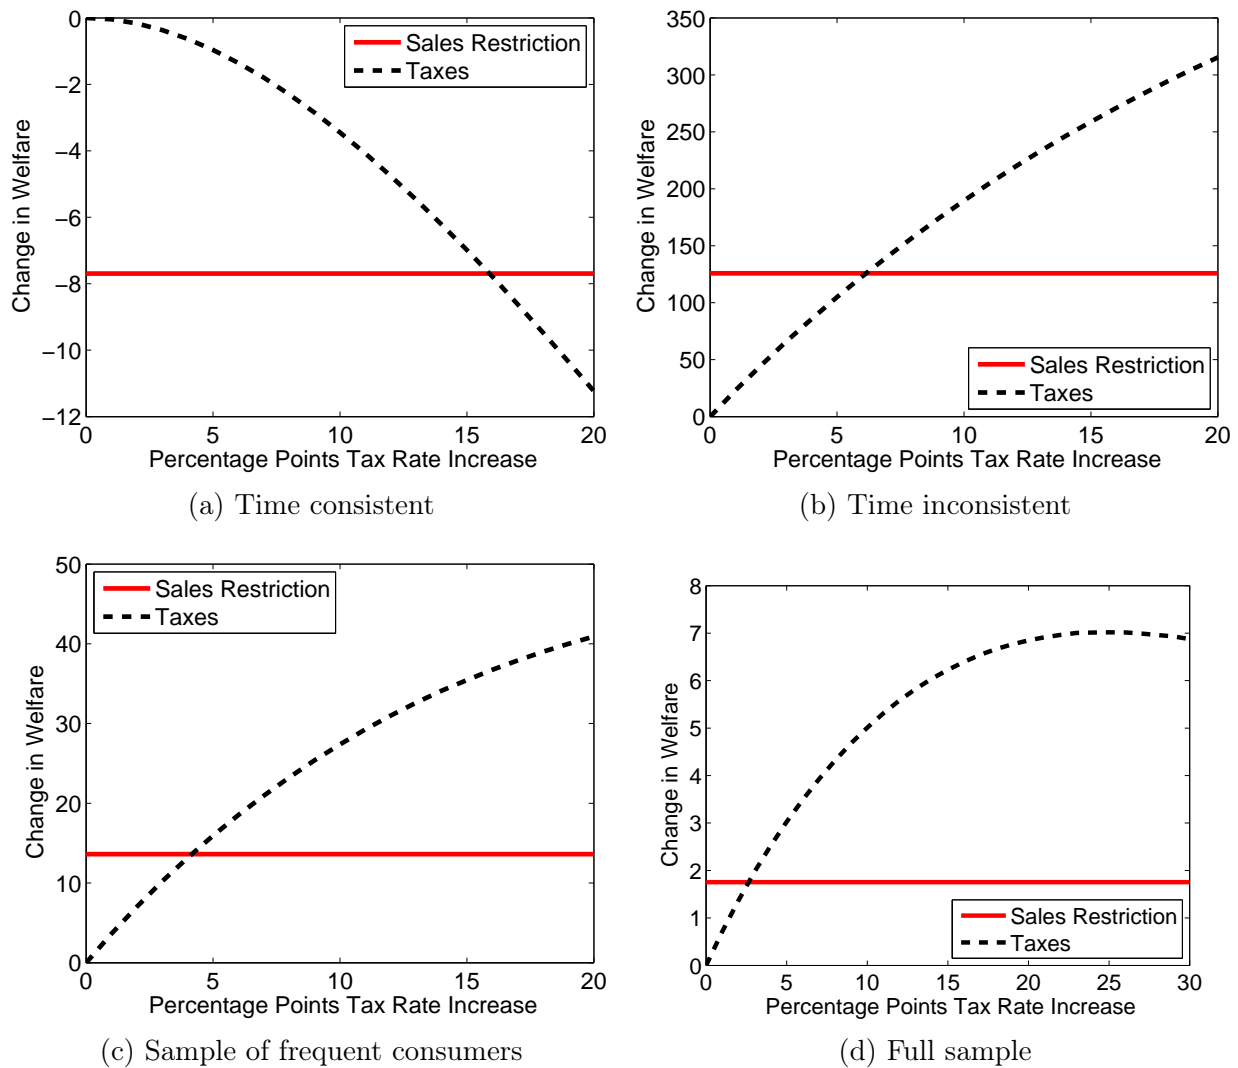


Figure 8: Change in consumer welfare due to the sales restriction and taxes

Note: Subfigures (a)-(c) use parameter estimates from column 1 in Table 6 and data from the sample of frequent consumers. Subfigure (d) uses parameter estimates from column 5 in Table 6 in addition to the above and data from the full sample. Time inconsistent consumer's welfare is his 'long term self's' welfare (utility function without time inconsistent discounting β). On y-axes, the change in welfare is measured in dollars and annualized.

A Procedures to Clean the Raw Data

The dataset is from the Nielsen Homescan Consumer Panel. It covers the time period from 2004 to 2014 and includes data from the following 10 states: Arizona, California, Florida, Illinois, Louisiana, Missouri, Nebraska, Nevada, South Dakota, and Wisconsin.

I include in the sample only purchases and shopping trips to the types of stores, where consumers commonly buy beer or groceries. For example, shopping trips to apparel stores won't be included, because in an apparel store consumers typically cannot buy beer. I also exclude online purchases, because of the gap between the time of the purchase and the arrival of the product. Specifically, I include the following types of stores: grocery stores, discount stores, liquor stores, drug stores, warehouse clubs, convenience stores, service stations, gas mini marts, military stores, and beverage stores. Purchases in these stores account for over 97% of all beer purchases in my dataset.

I exclude households who have ever bought beer that has missing product characteristics. Since my analysis mainly concentrates on weekends and I would like to see the behavior across the entire weekend, I will exclude the weeks when households are in the consumer panel only for a half weekend (in the beginning and end of panel years). After that I'm left with 27,419 households that form the full sample in my analysis.

For most of the analysis, I restrict attention to frequent consumers who purchase beer rather often. Specifically, I restrict the main sample in the following way. (1) I exclude a household in a given year if it purchased beer strictly less than in 10 weeks per year. This leaves me with 6,234 households. Households can be different over the years. If a household satisfies the requirement one year, but not on the others, then it is included in the sample only that year. (2) I also exclude households, who do not purchase beer on weekends. Specifically, who purchase beer in less than 5 weekends over the observable years. This leaves me with 5,385 households that form the sample of frequent consumers, while the remaining 22,034 form the sample of infrequent consumers.

I treat beer as a homogeneous good and I collapse the data into a household-day level. For each beer purchase, I observe each household purchasing a particular brand, package size and quantity at a specific time and at a specific store. I collapse the household beer purchases into one purchase per day, such that each observation is characterized by a particular household in a particular day and the total purchased quantity (in dollars and volume) is the sum of all his beer purchases on that day.

I calculate prices also from the Homescan Consumer panel. The price is calculated as a weighted average of the prices of all beer sales in the state in each week; prices are weighted

by sales.

B State with High Shopping Costs

Consider the state where shopping costs are high in both periods. The quantity that consumer purchases for period t consumption equals $\max\left\{0, \frac{1}{\beta(h+\alpha p)} - \eta_t\right\}$ if consumer buys it in period t , and it equals $\max\left\{0, \frac{1}{(h+\alpha p)} - \eta_t\right\}$ if consumer stores it as an inventory in period $t - 1$.

If consumer purchases something in period one, then he buys inventory, if $\eta_2 < \frac{1}{(h+\alpha p)}$. Let's now look at when consumer purchases something. Let's first consider the case, where discount factor is not too low. The consumption shock realizations, where consumer purchases in period one and two, are described on Figure 9 and given with the following equations.

$$Prob(x_1 > 0, x_2 = 0|CC) = \int_0^\infty \int_0^{\bar{\eta}_1(\eta_2)} dF_1(\eta_1)dF_2(\eta_2) \quad (17)$$

and

$$Prob(x_1 = 0, x_2 = 0|CC) = \int_{\frac{\Lambda(C)}{\beta(h+\alpha p)}}^\infty \int_{\bar{\eta}_1(\eta_2)}^\infty dF_1(\eta_1)dF_2(\eta_2) \quad (18)$$

$$Prob(x_1 = 0, x_2 > 0|CC) = \int_0^{\frac{\Lambda(C)}{\beta(h+\alpha p)}} \int_{\bar{\eta}_1(\eta_2)}^\infty dF_1(\eta_1)dF_2(\eta_2) \quad (19)$$

where

$$\bar{\eta}_1(\eta_2) = \begin{cases} \frac{\Lambda(C)}{\beta(h+\alpha p)} & \text{if } \frac{1}{(h+\alpha p)} < \eta_2 \\ \frac{\Lambda(C - \beta\Lambda^{-1}((h+\alpha p)\eta_2))}{\beta(h+\alpha p)} & \text{if } \frac{\Lambda(C)}{h+\alpha p} < \eta_2 < \frac{1}{(h+\alpha p)} \\ \frac{\Lambda((1-\beta)C)}{\beta(h+\alpha p)} & \text{if } \eta_2 < \frac{\Lambda(C)}{h+\alpha p} \end{cases}$$

When discount factor is low enough, then in addition to the above, demand in both periods may be positive as described on Figure 10.

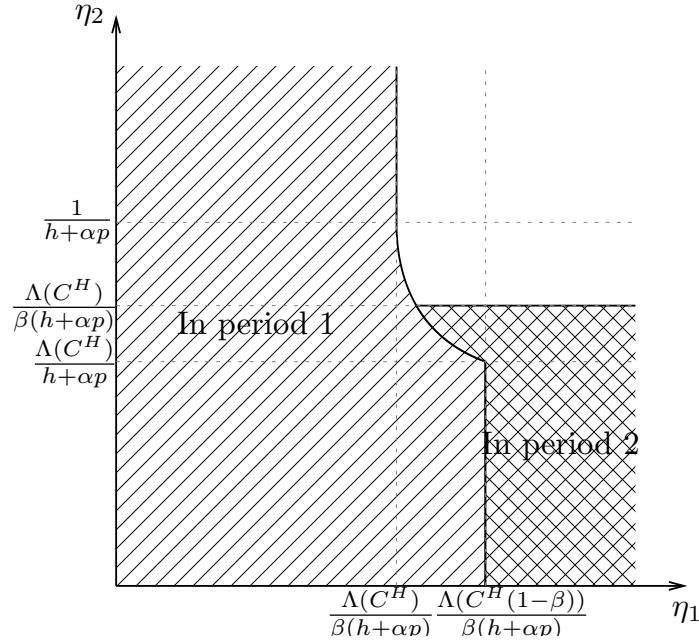


Figure 9: Probability of purchasing in period 1 and 2 when shopping costs are high both periods and discount factor is high.

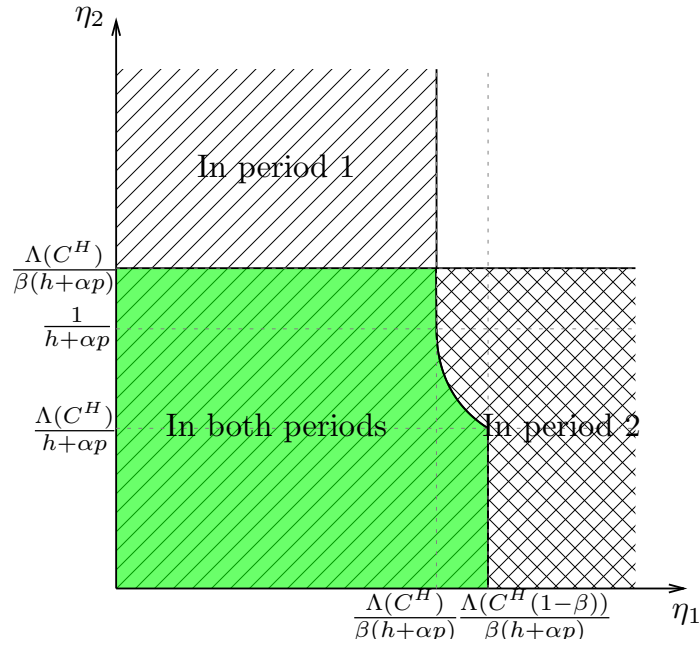


Figure 10: Probability of purchasing in period 1 and 2 when shopping costs are high both periods and discount factor is low.

C Contribution of Consumer i to the Log-Likelihood

The log-likelihood of observing the purchases of consumer i can be separated into the log-likelihoods of states: $\log L_i = \sum \log L_i^S$.

In state 00, I match only the probability of purchasing:

$$\begin{aligned} \log L_i^{00} &= \sum_{n=1}^N 1[00] \cdot \{1[x_{1n} = x_{2n} = 0] \cdot \log Prob(x_{1n} = x_{2n} = 0|00) \\ &\quad + (1 - 1[x_{1n} = x_{2n} = 0]) \cdot \log (1 - Prob(x_{1n} = x_{2n} = 0|00))\} \end{aligned} \quad (20)$$

where

$$\begin{aligned} Prob(x_{1n} = x_{2n} = 0|00) &= \left[1 - \Phi \left(\frac{-\log \beta(h + \alpha p_{1n}) - \mu_1}{\sigma_1} \right) \right] \\ &\quad \cdot \left[1 - \Phi \left(\frac{-\log \beta(h + \alpha p_{2n}) - \mu_2}{\sigma_2} \right) \right] \end{aligned} \quad (21)$$

In state C0, I match the probability of purchasing on Saturday and the quantity purchased on Sunday given no purchase on Saturday. I will not try to match the exact quantities purchased, but instead quantity intervals, defined by the following quantity thresholds $X_0 = 0 < X_1 < X_2 < X_3$, where X_1, X_2 , and X_3 are equivalent (in terms of total amount) to 6-pack, 12-pack, and 18-pack of 12 oz cans. The likelihood in the state can be written as:

$$\begin{aligned} \log L_i^{C0} &= \sum_{n=1}^N 1[C0] \{1[x_{1n} > 0] \cdot \log Prob(x_{1n} > 0|C0) \\ &\quad + 1[x_{1n} = 0, x_{2n} = 0] \cdot \log Prob(x_{1n} = 0, x_{2n} = 0|C0) \\ &\quad + \sum_{k=1}^3 1[x_{1n} = 0, X_{k-1} < x_{2n} \leq X_k] \cdot \log Prob(x_{1n} = 0, X_{k-1} < x_{2n} \leq X_k|C0) \\ &\quad + 1[x_{1n} = 0, x_{2n} > X_3] \cdot \log Prob(x_{1n} = 0, x_{2n} > X_3|C0)\} \end{aligned} \quad (22)$$

By assumption, in state C0, the first period purchases don't depend on the second period shock and if $x_1 = 0$, then the second period purchases don't depend on the first period shock. Therefore, the conditional probability of second period purchase equals the unconditional probability and the likelihood can be rewritten using the following probabilities.

$$Prob(x_{1n} = 0|C0) = 1 - \Phi \left(\frac{\log \Lambda(C_{1n}) - \log \beta(h + \alpha p_{1n}) - \mu_1}{\sigma_1} \right) \quad (23)$$

$$Prob(x_{2n} = 0|x_{1n} = 0, C0) = 1 - \Phi\left(\frac{-\log \beta(h + \alpha p_{2n}) - \mu_2}{\sigma_2}\right) \quad (24)$$

$$\begin{aligned} Prob(X_{k-1} < x_{2n} < X_k|x_{1n} = 0, C0) &= \Phi\left(\frac{\log \max\{\frac{1}{\beta(h+\alpha p_{2n})} - X_{k-1}, 0\} - \mu_2}{\sigma_2}\right) \\ &\quad - \Phi\left(\frac{\log \max\{\frac{1}{\beta(h+\alpha p_{2n})} - X_k, 0\} - \mu_2}{\sigma_2}\right) \end{aligned} \quad (25)$$

and

$$Prob(x_{2n} > X_k|x_{1n} = 0, C0) = \Phi\left(\frac{\log \max\{\frac{1}{\beta(h+\alpha p_{2n})} - X_k, 0\} - \mu_2}{\sigma_2}\right) \quad (26)$$

In state 0C, I match both the probability of purchasing on Saturday (when shopping is costless) and on Sunday (when shopping is costly).

$$\begin{aligned} \log L_i^{0C} &= \sum_{n=1}^N 1[0C] \cdot \{1[x_{1n} = 0] \cdot \log Prob(x_{1n} = 0|0C) \\ &\quad + 1[x_{1n} > 0, x_{2n} > 0] \cdot \log Prob(x_{1n} > 0, x_{2n} > 0|0C) \\ &\quad + 1[x_{1n} > 0, x_{2n} = 0] \cdot \log Prob(x_{1n} > 0, x_{2n} = 0|0C)\} \end{aligned} \quad (27)$$

In the last state, I match only the probability of purchasing in both periods:

$$\begin{aligned} \log L_i^{CC} &= \sum_{n=1}^N 1[CC] \cdot \{1[x_{1n} > 0, x_{2n} > 0] \cdot \log Prob(x_{1n} > 0, x_{2n} > 0|CC) \\ &\quad + (1 - 1[x_{1n} > 0, x_{2n} > 0]) \cdot \log (1 - Prob(x_{1n} > 0, x_{2n} > 0|CC))\} \end{aligned} \quad (28)$$

D Estimates with Alternative Datasets

Table 9: Parameter estimates with alternative datasets

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ρ	0.845 (0.012)	0.841 (0.006)	0.796 (0.022)	0.741 (0.102)	0.904 (0.029)	0.741 (0.015)	0.885 (0.013)
β	0.459 (0.025)	0.598 (0.021)	0.481 (0.035)	0.465 (0.078)	0.611 (0.097)	0.524 (0.017)	0.334 (0.029)
C^H	0.237 (0.033)	0.112 (0.036)	0.213 (0.037)	0.413 (0.082)	0.104 (0.111)	0.170 (0.016)	0.431 (0.054)
α	0.042 (0.002)	0.038 (0.001)	0.041 (0.002)	0.050 (0.004)	0.087 (0.002)	0.042 (0.001)	0.042 (0.001)

Note: Each column presents estimates of the same model as in Column 1 in Table 6 using different datasets. In all cases the sample is restricted to the households who frequently buy beer. The sample in Column 1 excludes holidays. The sample in Column 2 excludes weeks with previous beer purchases during the week days. In Column 3, the sample includes only single person households. In Column 4, the sample includes only households who always purchase beer at either grocery stores or discount stores. In Column 5, the sample includes only the households who purchase only smaller quantities of beer (up to 18 12-ounce bottles/cans per day). Datasets in Column 6 and 7 use alternative threshold for shopping cost: the thresholds equal one third and two thirds, respectively. Bootstrapped standard errors are included in paranthesis.