Consumer-Financed Fiscal Stimulus: Evidence from Digital Coupons in China*

Jing Ding, Harbin Institute of Technology

Lei Jiang, Tsinghua University

Lucy Msall, University of Chicago

Matthew J. Notowidigdo, University of Chicago and NBER

First Version: February 25, 2022

This Version: July 21, 2023

Abstract

Beginning in 2020, local governments in China issued digital coupons to stimulate spending in targeted categories such as restaurants and supermarkets. We study this new form of fiscal stimulus using detailed data from a large e-commerce platform and a bunching estimation approach, and we find that the coupons caused large increases in spending of 3.5-3.6 yuan per yuan spent by the government. We find no evidence that the large spending effects come from substantial substitution away from non-targeted spending categories or short-run intertemporal substitution. To rationalize these results, we develop a dynamic consumption model to show how the coupon's minimum spending thresholds (i.e., "spend at least ¥X, get ¥Y off") create temporary notches that lead to large behavioral responses. In the model, the increased spending caused by the coupons is partially "consumer-financed", since consumers use their own money to push their spending above the coupon thresholds. Calibrating the model to match our empirical results, we find that coupons generate about half of the increase in consumer welfare as an equivalent amount of government spending distributed as cash, but that coupons are a much more cost-effective form of targeted fiscal stimulus.

*We thank I-Han Lai, Jeremy Mopsick, and Ningxin Zhang for excellent research assistance, and we thank Milena Almagro, Emir Kamenica, Koichiro Ito, Zhiguo He, Nancy Qian, Peter Ganong, Anders Humlum, Kory Kroft, Eric Zou, Eric Zwick, and numerous seminar participants for helpful comments. We gratefully acknowledge funding from the Tsinghua University - University of Chicago Joint Research Center for Economics and Finance. Msall acknowledges support from the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1746045. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Gao Yuan (Interviewer): *Some developed countries have opted for cash. Why do you think China should issue consumer coupons as the main means of stimulus?*

Justin Yifu Lin (World Bank Chief Economist, 2008-2012): *The situation in China is different. If cash is distributed, except for a few disadvantaged groups who will immediately go to buy necessities, most people will probably deposit the money in the bank and not necessarily consume it. It is difficult to achieve the dual function of protecting the family and protecting the enterprise.*

Jiefang Daily, May 31, 2020

---

1 Introduction

Countries all around the world distribute stimulus payments during economic downturns to increase consumption. In the United States, for example, the federal government distributed billions of dollars of stimulus payments to households in each of the last three recessions, and each time households used the payments to immediately increase their consumption (Johnson et al. 2006; Shapiro and Slemrod 2009; Parker et al. 2022).

Many governments also design fiscal stimulus policies to target particular sectors of the economy. For example, during the 2008-2009 Great Recession the US government provided targeted financial support for the automobile industry through the "cash for clunkers" program and supported the real estate industry through a new first-time homebuyer tax credit (Mian and Sufi 2012; Berger et al. 2020).

More recently, during the 2020-2021 COVID-19 recession, provinces and municipalities across China carried out a novel form of targeted economic stimulus using government-issued digital coupons. The coupons were delivered online through smartphone apps and designed to encourage spending in certain categories such as restaurants, food delivery, grocery stores, and entertainment. In China (like in many other parts of the world), these sectors were hit particularly hard during the early months of the global COVID-19 pandemic and many provincial and municipal governments in China designed these coupons to stimulate spending in these sectors as their local economies re-opened following the initial wave of the virus.

In addition to the digital distribution, another major difference between the coupons and other forms of stimulus is that the coupons had fixed spending thresholds that needed to be reached before consumers received money from the government – for example, one coupon would give 18 yuan off of a food delivery order if the total transaction amount was at least 54 yuan ("Spend at least ¥54, get

¥18 off").

In this paper, we estimate the effects of the digital coupons on consumer spending, and we evaluate the coupons' effectiveness as fiscal stimulus. To do this, we assemble data covering several different types of coupons distributed across three cities in China. The different coupons have a range of different thresholds and apply to several different spending categories. Throughout this paper, we define the "coupon MPC" to be the increase in consumption caused by the coupons relative to the fiscal cost of the coupons. For example, if 200,000 "Spend at least ¥54, get ¥18 off" food delivery coupons were distributed in a city, and 100,000 of them were redeemed, then the direct fiscal cost is $18 \times ¥100,000 = ¥1,800,000$. If the total increase in spending caused by the coupons is $¥4,500,000$, then we conclude that $MPC^{coupon} = 2.5$.

The reason that the coupon MPC can be larger than one is that many consumers may need to increase their spending substantially in the targeted spending category in order to reach the threshold and be able to take advantage of the coupon. If they do not decrease their spending in other categories (which would "offset" the increase in the "targeted" spending category), then total consumer spending would increase by more than the discount associated with the coupon, which is the amount financed by the local government. As a result, we call this new form of fiscal stimulus *consumer-financed fiscal stimulus*, since the increased spending is caused by government spending but is ultimately mostly paid by consumers.

Our empirical analysis uses detailed account-level and transaction-level data from a large online shopping platform that distributed the coupons across many cities and provinces in China. The data set contains detailed information about each transaction including the spending amount and the time and date of each transaction. The account-level data also contains information on everyone who received coupons and whether or not the coupons were redeemed during the time window when the coupons were valid.

We begin our analysis by presenting clear visual evidence of sharp "bunching" around coupon-specific thresholds during the weeks that the coupons could be used. We find no evidence of similar bunching in the weeks before and after the coupons were distributed. These results indicate a clear behavioral response to the coupon-specific spending thresholds.

We then build on these graphical results by developing a bunching estimator that compares the number and distribution of transactions before, during, and after the weeks when the coupons were available to be used. Under the assumption that the distribution of transaction-level spending in the weeks before the coupons were distributed represents a valid counterfactual estimate of the distribution of spending in the weeks during and after the coupons were distributed, we can identify the effects of the coupons on spending by "differencing" the spending distributions and then integrating across the difference in spending distributions to recover the coupon MPC for each coupon. This empirical approach is broadly related to previous work in public economics and labor economics which uses a similar kind of "bunching estimator" to infer the behavioral responses to tax kinks, tax notches, and minimum wages (Best et al. 2020; Defusco et al. 2020; Cengiz et al. 2019). In our setting, the coupons create a notch in the consumer's budget constraint, similar to the recent Kleven and Waseem (2013) analysis of tax notches in Pakistan.

Turning to our main results, we find coupon MPC estimates ranging from 1.1 to 4.1, with a weighted average of 3.3-3.6. These estimates are stable for at least several months after the coupons are distributed. We assess whether the large coupon MPCs we estimate come partly from substitution between "targeted" and "non-targeted" spending categories, and we find no evidence of meaningful cross-category substitution, suggesting that the coupons increased aggregate consumer spending in the short run. We also find very little intertemporal substitution in the short run, finding very large MPCs even going out several months after the coupons were distributed. This implies that simple re-timing of consumption in the short run is not likely to account for much of the large coupon MPCs we estimate.[1]

Our main results are robust to several alternative ways of implementing the bunching estimator, and we are also able to produce estimates from an alternative empirical approach that exploits the explicit random assignment of coupons for a subset of the coupons in our data. In these cases, consumers were randomly assigned one of three different coupons, which allows us to calculate the causal effect on spending from being assigned a high (or medium) threshold coupon compared to a low threshold coupon. We show that the causal spending effects from the random assignment are related to the bunching estimator estimates through an identity. As far as we know, this is the first time local extrapolation using a bunching estimator is assessed using explicit random assignment of the notch or kink that is used for identification in the bunching estimation.

In the final part of the paper, we develop a simple dynamic model of consumer spending to understand the economics behind our reduced-form results. We use the model to derive formulas for the coupon MPC and the effect of coupons on consumer welfare relative to the equivalent amount in cash. We are not aware of any existing estimates of the MPC out of cash in China, but we expect it to be similar – or perhaps slightly smaller than – the recent estimates in the U.S. during the same time period (see, e.g., Parker et al. 2022). We calibrate our simple model to match the very large coupon

[1]This contrasts with the findings in the Mian and Sufi (2012) analysis of the "cash for clunkers" program, where the initial short-run increase in auto purchases from the stimulus program comes almost entirely from the re-timing of auto purchases rather than from an aggregate (net) increase in auto spending during the year of the policy.

MPCs we estimate, and show that it can still match a small MPC out of cash. The key to matching our reduced-form results is that the coupon threshold needs to be set higher than the spending many consumers would have wanted to choose in the targeted sector in the absence of the coupon, which appears to be a reasonable assumption given the location of the threshold in the pre-period spending distribution.

We also use the calibrated model to illustrate how the coupon MPC varies with the coupon's threshold and discount amount and to calculate the welfare cost to consumers from receiving a coupon instead of cash. We estimate that consumers get about 50 percent of the increase in consumer welfare they would get from equivalent amount of fiscal stimulus distributed as cash, but the targeted sector receives much more spending from coupons compared to cash, highlighting the potentially attractive targeting properties of this novel form of stimulus. Lastly, we use the calibrated model to simulate alternative coupon designs find that lower coupon thresholds and higher discount amounts would be less cost-effective but would deliver greater aggregate stimulus.

Taken together, our empirical and theoretical results suggest that the digital coupons may be a particularly cost-effective way to provide targeted fiscal stimulus to specific sectors. The coupon MPCs that we estimate are large and persistent for several weeks, implying that the increased spending is achieved at a very low fiscal cost compared to other forms of stimulus. As a result, we conclude that the digital coupons represent a technological innovation in stimulus policy that may be particularly attractive when there are specific sectors that need targeted support during economic downturns.

When it comes to fiscal stimulus, at least, the highly non-linear incentives created from the notch-based coupon design may be a feature rather than a bug. Just as tax notches sometimes create unusually large behavioral responses by taxpayers, the digital coupons we study in this paper consistently create unusually large spending responses because the thresholds temporarily create notches in the household's budget constraint. While tax notches are usually seen as a "design flaw" in public finance (since it's often difficult to imagine optimal tax policy featuring notch), in our setting the coupon-specific spending thresholds create large behavioral responses in a way that may deliver cost-effective stimulus to targeted sectors compared to both direct cash payments to households and direct subsidies to the targeted industry.

Our paper contributes to three main areas of research. First, the paper contributes to a large and active literature on the consumption responses to different types of fiscal stimulus. Much of this literature in recent years has focused on spending responses to direct cash payments, but there are also papers studying other forms of fiscal stimulus such as the "cash for clunkers" and first-time homebuyer credit policies mentioned above, as well as recent related work studying shopping coupons in Japan and shopping vouchers in Taiwan (Kan et al. 2017; Hsieh et al. 2010). Relative to these studies, our focus is on novel digital coupons during the COVID-19 recession that had coupon-specific spending thresholds, and we develop a model to explain why these thresholds appear to be primarily responsible for the particularly large coupon MPCs compared to these other policies.

Second, this paper is related to recent papers studying behavioral responses to "notches". Tax notches create strong incentives for behavioral responses to avoid dominated choices. The same kind of incentives are created by the digital coupons we study in this paper. For households receiving the coupons, as long as there is "free disposal" there is no reason for households to spend just below the coupon threshold when a coupon is available. Unlike in Kleven and Waseem (2013), where the dominated choices are used to reveal the extent of optimization frictions, in our setting the share of households with coupons who make dominated choices in the same way is extremely small. Some of this may be due to the fact that households had a choice about whether to try to get a coupon, and for the households who chose to try to receive the coupons, they very rarely ended up choosing transaction amounts just below the threshold needed to unlock the discount. This is arguably another benefit of the coupon stimulus policy, which is that the coupons are only "taken up" by households who are interested and able to increase spending enough to take advantage of the threshold-based benefit.

Lastly, this paper is most closely related to two other recent studies of digital coupons using different data sets and different empirical approaches. Xing et al. (2021) study digital coupons in a single large Chinese city and find an average MPC of around 3.0 comparing consumers who just barely missed out on getting a coupon ("near miss") to consumers who just barely received a coupon. These estimates are broadly similar to our own main results, and we see the two projects as highly complementary since the papers have different methodologies and identifying assumptions, but similar coupon MPC estimates.[2] Liu et al. (2021) use administrative data on coupons issued on Alibaba in Hangzhou and Guangxi and find that the marginal propensity to consume (MPC) to the range of 3.4-5.8. This paper takes a difference-in-difference approach comparing consumers who received coupons to a random sample of individuals who tried but failed to get a coupon. The approach is similar in spirit to the Xing et al. (2021) approach, although it is not set up as a "near-miss" design, so it requires somewhat stronger identifying assumptions, and still implicitly relies on there being excess

[2] One benefit of the "near miss" approach is that it also rules out anticipatory behavior as a source of the large coupon MPCs. The fact that our results are similar to the results in Xing et al. (2021) suggests that anticipatory behavior is not a major reason for the large MPCs in our setting. Xing et al. (2021) also studies how the coupons cause consumers to shift consumption between firms and find that the coupons cause consumers to spend more at larger firms that sell pricier goods and services.

demand for coupons so that the limited offering of coupons is "binding."

Relative to both of these Chinese coupon papers, our analysis covers a larger number of cities and studies more coupons. Our bunching estimator approach is also a distinct research design that is not used in either of the two previous papers. We are also the first study we are aware of that estimates coupon MPCs using explicit random assignment of coupons, which provides validation of both our bunching estimates and the other estimates in the literature. Lastly, we develop and calibrate a model that can simultaneously fit the large estimated coupon MPCs alongside a small MPC out of cash, which we use to measure the consumer surplus from coupons relative to cash and to evaluate counterfactual coupon designs.

The next section provides background on the digital coupons that we study. Section 3 describes our data and sample construction. Section 4 describes our empirical approach and formally defines the coupon MPC concept we use throughout the paper. Section 5 describes our main results, assesses the importance of intertemporal substitution and substitution to other spending categories, and validates our results using random assignment of coupons. Section 6 presents a simple model of spending responses to the digital coupons and cash and calibrates the model for welfare analysis and counterfactual scenarios. Section 7 concludes.

2 Background on the Chinese Coupon Programs

The Covid-19 pandemic at the beginning of 2020 dramatically slowed China's economy. In response to this unexpected macroeconomic shock, provincial and municipal governments in China issued short-term Internet-based digital coupons to stimulate the economy from the demand side[3]. The digital coupons were distributed through pre-existing consumer spending technology platforms such as Alibaba, Meituan, and JingDong. The coupons were distributed directly to consumers through multiple waves in many cities across China. The first city to issue coupons was Jinan, the capital of Shandong Province, whose government offered a total of ¥20 million worth of coupons on March 2, 2020.

The stated aim of the stimulus policy was to promote consumption of local individuals and enhance the demand of the local economy at relatively low fiscal cost in response to the Covid-19 pandemic and the subsequent economic recession. Different types of coupons could only be used in their specific categories in order to help the recovery of those industries hit hardest by the pandemic (such as restaurants and tourism). The central government in China only gave general guidance to the local

[3]Other demand-side Chinese stimulus policies during this time period included shopping festivals and subsidies for purchasing durables. On the supply side, Chinese policies included tax cuts and fee cuts for firms and the provisions of preferential loans to firms to help the firms resume production in the so-called "street-stall" and "night-time" economy.

governments that they ought to undertake these policies, which meant that the final coupons that were released by local governments differed across locations in terms of important features such as issuing volume and the types of consumption targeted. However, most of the issued coupons were in one of the following seven categories: food, retail, cultural tourism, automobiles, home appliances, sports, and information. Additionally, all of the coupons had thresholds and discount amounts – i.e., "Spend at least ¥X, get ¥Y off" – and all of the coupons had a fixed, short time period when the coupons needed to be used before they expired ("use it or lose it"). Although the coupon programs began during the worst of the Covid crisis, municipalities across China continued to offer additional coupons in 2021, 2022, and 2023.

We now describe the coupons that we analyze in more detail. Our data cover coupons distributed in three different cities during 2021 and cover a range of spending categories. We anonymize the cities in order to help protect the anonymity of the platform that provided the data for our study. Throughout this paper we describe coupons with a spending threshold X and a discount Y as "X-Y" coupons.

2.1 City A Coupons

City A is the capital city of a province, with a population of over 12 million. According to Yicaid Global, one of the major Chinese financial media outlets, it is a "new first-tier" city[4]. GDP in 2020 was ¥1.56 trillion, although this number reflects a substantial negative shock from the Covid-19 pandemic. The per capita disposable income of urban residents is approximately ¥50,400.

In February 2021, the City A municipal government announced that digital coupons would be rolled out during the Spring Festival through multiple mobile platforms, with the objective to stimulate the local economic recovery after Covid-19 pandemic and to reduce residents' outbound travels during the Spring Festival.

All residents in City A were eligible to take up one coupon through the mobile platform in each "wave" of coupon issuance. The number of coupons in each wave was limited. Residents in City A had to acquire the coupons from the online platform on a "first-come-first-service" basis. We evaluate the first two waves of coupons in City A. The first-wave coupons were released on February 8 and were valid for redemption between February 8 and February 16; the second-wave coupons were issued

[4] Since 2017, Yicai Global has categorized Chinese cities into six tiers – First-tier, New first-tier, Second-tier, Third-tier, Fourth-tier, and Fifth-tier. Cities rankings reflect differences in consumer behavior, income level, population size, infrastructure readiness, transportation network, and business opportunity. The cities we evaluate (Cities A, B, C) were classified as New first-tier, Third-tier, and Second-tier respectively, each year from 2017 to 2022. Source from: http://www.datayicai.com.

on February 17 and expired on February 26.

Coupons available on our platform during these waves included Supermarket coupons (which could be redeemed only in online supermarkets), "Life Service" coupons which could be redeemed in food delivery, movies, hotel booking, entertainment, or restaurants, and Shopping coupons which could be redeemed in physical shopping malls, hotels, and for other types of entertainment. We focus on the supermarket and life service coupons because of implementation issues with the shopping mall coupons.[5]

In each wave, Supermarket coupons were offered at three levels: "24-8", "54-18", and "84-28", where the first number refers to the spending threshold, and the second number refers to the discount the coupon granted if the threshold was reached. Similarly, three levels of life service coupons were available: "54-18", "84-28", and"114-38". Combining across supermarket and life service coupons, City A offered ¥19 million worth of coupons on our mobile platform during these two waves.

A unique feature of City A's coupons is that when residents took up a supermarket or life service coupon, they could not choose between the three possible levels. The threshold and discount was randomly assigned by the platform between the three possible threshold-discount pairs. In particular, each coupon recipient opted in to receive a coupon (of a particular spending category) without knowing the threshold and discount of the coupon they were getting; the residents only knew it would be one of the three types listed above.

2.2 City B Coupons

City B is a prefecture-level city in the south-east of China – a "third-tier" city according to Yicai Global's classification. As of the 2020 census, its population was 3.37 million inhabitants. The city's GDP was ¥320 billion in 2020, with the per capita disposable income of urban residents approaching ¥62,000.

The City B municipal government offered ¥2.89 million worth of food delivery coupons on our platform in 2021, in two separate waves. In the first wave, distributed between July 30 and August 6, 2021, all coupon recipients received a bundle of two identical "30-15" food delivery coupons which were valid for 15 days. The coupons were redeemable in online food delivery purchases in City B when

[5] We exclude the shopping coupons because retailers in this category created store-specific or shopping-mall-specific coupons which they allowed holders of government coupons to purchase, thus effecting an exchange of a quick-expiring government coupon for a longer-duration private coupon. This was controversial behavior since it is a kind of "duration arbitrage", but it also makes it impossible to reliably measure the timing of actual spending. We only see the "exchange" of the government coupons into retail coupons in our data.

a purchase met the minimum requirement of ¥30.

In the second coupon wave, from August 26 to September 2, recipients could choose one or both of two available coupons: a "30-15" coupon and a "20-10" coupon. Each coupon was redeemable for online food delivery purchases for three days from the time of receipt. Residents are allowed to claim up to two coupons (a maximum of one of each type) every three days during the coupon wave.

2.3 City C Coupons

City C is a coastal prefecture-level city in the east of China – a "second-tier" city according to Yicai Global's classification. The total population in 2020 was 7.1 million. The per capita disposable income of urban residents stood at ¥49,400, slightly higher than the national average (¥43,800). The gross domestic product of City C in 2020 amounted to approximately ¥ 781 billion.

In the last quarter of 2021, the City C municipal government offered ¥3.3 million worth of digital coupons in two waves, from October 20 to October 29 and from October 30 to November 8, respectively. Each coupon wave included two types of coupons: 100-40 and 200-100. Residents could choose only one type of coupon. Each was redeemable in the "life service" categories of food delivery, hotel booking, or restaurants. Take-up of coupons in City C was particularly high. For example, in the second release wave, it took only 4 seconds before all the "200-100" coupons were claimed, and it took only 16 seconds before all "100-40" coupons were claimed.

2.4 Comparison to Other COVID-era Coupons Outside of China

Several countries outside China offered coupon programs to promote consumer spending during the pandemic. However, these "coupons" did not feature a threshold design.

Beginning in December 2020, Italy's "Cashless Italia" program offered a 10% rebate, up to €150 per half-year for non-business transactions made via digital payment in physical stores. Participants did not need to meet a minimum spending threshold, but individuals had to register their cards with the government and make at least 50 transactions on their card during the half year to participate. The primary motivation of the programs seems to have been to move Italy towards a cashless economy, but a secondary motivation may have been to support brick-and-mortar businesses over online retailers. The program was suspended in late June of 2021 (Italian Ministry of Economy and Finance 2021).

The United Kingdom attempted to promote consumer spending in the hospitality sector through it's "Eat Out to Help Out" program. On Mondays, Tuesdays and Wednesdays during the month of August 2020, the government subsidized 50% of food and non-alcoholic drink purchases at participating restaurants, up to a total of £10 per person per visit (Government of the United Kingdom 2020).

We are not aware of any other coupon programs with a threshold design like the ones used in the digital coupons in China.

3 Data

We use detailed transaction data from one of the large Chinese shopping platforms that distributed the coupons. This e-commerce platform (henceforth, "the platform") has had substantial market share for spending categories including entertainment, dining, and food delivery for several years. By 2018, the platform had more than 600 million registered users with more than 290 million of them actively using its service. The platform has roughly 35 million daily users. Our data cover coupons issued in the three cities described in the previous section.[6]

For each transaction in the sample, we observe the consumer's de-identified account number, transaction number, transaction time, transaction amount, and transaction spending category. We have access to transactions for three months before and after the coupons were distributed.

We merge the transactions data with the coupon database, which records each coupon holder's de-identified account number, the time the coupon was acquired, the coupon face value, and the coupon threshold. The coupon database identifies each coupon holder in each coupon wave. We also observe coupon redemption information in a separate coupon redemption database. For each redeemed coupon, we observe the transaction number, transaction amount, coupon id, discount amount, and the actual payment.

By merging the transactions data with the coupon database using the account numbers, we are able to get all of the transactions for all of the coupon holders each coupon wave, whether or not the coupons were actually used. This level of detail is important because the coupon redemption rate is often substantially below 100 percent for all of the coupons in our data. Most of our analysis proceeds using the distribution of spending at the transaction level, focusing on the full set of consumers who acquired coupons each coupon wave.

Table 1 presents summary statistics for the coupons in our data and reports the total coupons available, the coupons taken up, and the coupons redeemed. Table 2 converts these statistics into take-up rates and redemption rates. There is variation in both the take-up rate and redemption rate

[6]The data was provided under a data use agreement which requires us to preserve the anonymity of the platform and the analyzed cities. The platform reviewed the study prior to public dissemination, but only for factual inaccuracies, confidential information, and trade secrets. All of the researchers had access to the "binned" transaction data that forms the basis of the bunching estimator described in more detail below.

across coupons. The redemption rate is always below 50 percent, which means that most coupons taken up by consumers are not used.

Table 2 shows that the take-up rate for City A coupons is constant within wave and spending category. This is consistent with strict random assignment of the coupons within a coupon category and wave in this subset of coupons. The variation in redemption rate is consistent with a causal impact of the coupon threshold and discount amount on the likelihood that the consumer chooses to use the coupon during the time window.

Figure 1 shows the range of coupon thresholds and discounts in our data. Almost all of the coupons have fairly low thresholds (below ¥120), and the discounts are always set between 25 and 50 percent of the coupon threshold. In other words, when cities offered coupons with higher thresholds, they tended to choose higher discounts, as well.[7]

To create our final analysis data set, we define the period length for each coupon to be the coupon's period length (i.e., the number of days each consumer had to use the coupon before it expired), and we make sure to include the same days of the week as the coupon period to account for any possible day-of-week effects. For example, if a coupon was available to use for 5 days from Tuesday to Saturday, then we define our first pre-period to be the Tuesday to Saturday in the prior week. For some cities that experience multiple coupon waves in a row, we will take the pre-period prior to the first wave as the pre-period for all waves, in order to avoid overlap between any coupon treatment and the pre-period. When we examine spending during post-periods (i.e., the periods after the coupon wave), we define the post-periods similarly.

4 Empirical Approach

We begin by describing our bunching estimator, which is our main empirical approach. We then describe how we translate our bunching estimates into a coupon MPC estimate. Lastly, we describe how we use the random assignment of a subset of the coupons in our data to validate the bunching estimates.

[7]In interviews with employees of the platform, they described the municipalities as trying to hit a certain "leverage ratio" which they defined as the ratio of the coupon threshold to the coupon discount amount. This ratio is similar but not quite the same as the closed-form expression for the MPC^{coupon} that we derive in Section 6 below. See the Online Appendix for more institutional details we learned from these interviews.

4.1 Bunching Estimator

To estimate the effects of the coupons on spending, we use a modified bunching estimator that takes the distribution of spending in the periods before the coupons are distributed to be a plausible counterfactual distribution of spending in the absence of coupons. This approach follows Best et al. (2020), Defusco et al. (2020), and Cengiz et al. (2019) by using pre-treatment distributions to form the counterfactual. This approach avoids relying on strong parametric functional form assumptions that are needed to estimate behavioral responses in a single cross-section (see, e.g., Saez 2010; Chetty et al. 2011; Kleven and Waseem 2013; Kleven 2016; Blomquist et al. 2021). Specifically, we take the distribution of purchases for coupon-receivers during the pre-period immediately prior to coupon recipient to take advantage of the fact that our dataset tracks transactions of each individual account over time.

Given this setup, our bunching estimator takes as an input the distribution of spending in various time periods grouped into ¥1 bins. We then calculate the effects of the coupons on spending by calculating the "excess mass" of transactions above the coupon threshold and the "missing mass" of transactions below the coupon threshold. We do this by comparing the number of transactions in different bins during the coupon wave and the pre-period. This results in the following bunching estimators:

$$\hat{B}_\tau = \sum_{j=\tau}^{\tau+H} (n_j^{WAVE} - n_j^{PRE})$$

$$\hat{M}_\tau = \sum_{j=1}^{\tau-1} (n_j^{WAVE} - n_j^{PRE})$$

where $\hat{B}_\tau$ is the "excess mass" of transactions above the coupon's spending threshold, and $\hat{M}_\tau$ is the "missing mass" of transactions below the coupon's spending threshold. The parameter τ denotes the coupon-specific spending threshold, n_j^{PRE} is the number of transactions with spending amounts between j and $j+1$ yuan in the pre-period, and n_j^{WAVE} is the number of transactions with spending amounts between j and $j+1$ observed during the coupon wave period. The parameter H is a tuning parameter that defines the upper limit of the bunching window. Effectively, this parameter rules out the coupons having any effects on the spending distribution at spending levels greater than $\tau + H$.[8] We choose $H =$ ¥50 for all of the coupons in our data, but all of our main results are very robust to alternative upper limits (see Online Appendix).

[8]Intuitively, for consumers who would have spent amounts greater than H in the absence of a coupon, the coupon represents a pure cash transfer, with a "face value" given by the coupon's discount. If there were large effects in the upper tail of spending distribution, we would interpret this as a likely violation of the research design under the assumption that any income effects are trivially small.

The excess mass estimate $\hat{B}_\tau$ contains a mix of intensive and extensive margin responses. Some of the additional transactions contributing to the excess mass represent spending that would not have occurred at all in the absence of a coupon (extensive margin), while some transactions contributing to this excess mass represent intensive margin responses, because consumers decided to spend more than they otherwise would have in order to be able to use the coupon (intensive margin). The sum of the excess mass and missing mass estimates, $\hat{B}_\tau + \hat{M}_\tau$, is thus the total effect of the coupons on the total number of transactions.

4.2 Estimating MPC^{coupon}

The previous subsection described how to estimate the effect of coupons on the total number of transactions. This approach leads to a simple visual presentation of sharp bunching of transactions at the coupon threshold and tests the null hypothesis of no behavioral response to the coupons.

To quantify the magnitude of the behavioral response, we calculate additional bunching estimates that give the total increase in spending caused by the coupons. Instead of calculating the effects of coupons on the number of transactions, we calculate the effects of the coupons on total spending in the spending category targeted by the coupons. This leads to two additional bunching estimators:

$$\hat{\mathfrak{B}}_\tau = \sum_{j=\tau}^{\tau+H} (n_j^{WAVE} - n_j^{PRE}) \cdot j$$

$$\hat{\mathfrak{M}}_\tau = \sum_{j=1}^{\tau-1} (n_j^{WAVE} - n_j^{PRE}) \cdot j$$

where the net increase in spending is given by $\hat{\mathfrak{B}}_\tau + \hat{\mathfrak{M}}_\tau$.

We define the coupon MPC (MPC_τ^{coupon}) to be the net increase in spending divided by the total amount of spending by the municipal government on the coupons:

$$MPC_\tau^{coupon} = \frac{\hat{\mathfrak{B}}_\tau + \hat{\mathfrak{M}}_\tau}{S_\tau} \tag{1}$$

where S_τ to denote the total amount of subsidy spending by the government on coupons with threshold τ, which equals the per-coupon subsidy d times the total number of coupons that were redeemed during the coupon wave.[9]

[9]For example, the 24-8 coupon has a threshold $\tau = 24$ and a per-coupon subsidy $d = 8$. Table 1 reports that 20,846 of the 24-8 coupons were redeemed. Therefore, the total government subsidy for this coupon type was $S_{\tau=24} = 8 \times 20,846 =$ ¥$166,768$ RMB. See Appendix Table OA.4 for the subsidy calculations for the other coupons.

4.3 Bunching Estimates Using Random Assignment

In a subset of the coupons (those of City A), the coupons are randomly assigned within a spending category. This means that conditional on acquiring a coupon, the threshold and discount is chosen randomly from a set of three options. The residents did not know which coupon they would receive. As a result, we can estimate the causal effect of being assigned one coupon relative to another coupon by simply comparing the distribution of spending in the coupon wave. This leads to a different $MPC^{coupon}_{\tau-\tau'}$ estimate based on a comparison of spending of residents assigned the coupon with threshold τ to the residents randomly assigned the coupon with threshold τ'. We construct this alternative coupon MPC estimator as follows:

$$MPC^{coupon}_{\tau-\tau'} = \frac{\sum_{j=1}^{H}(n^{WAVE}_{j,\tau} - n^{WAVE}_{j,\tau'}) \cdot j}{S_\tau - S'_\tau} \tag{2}$$

where $n^{WAVE}_{j,\tau}$ is the number of transactions with spending amount between j and $j+1$ for the residents randomly assigned coupon with threshold τ.

We show in the Online Appendix that the coupon-specific bunching estimates MPC^{coupon}_{τ} are related by the following identity:

$$MPC^{coupon}_{\tau-\tau'} = \frac{S_\tau}{S_\tau - S'_\tau} MPC^{coupon}_{\tau} - \frac{S'_\tau}{S_\tau - S'_\tau} MPC^{coupon}_{\tau'} \tag{3}$$

In words, the identity above says that the MPC between any two pairs of randomly assigned coupons can be decomposed into a particular weighted average of the coupon MPCs estimated using bunching estimators. Suppose that both of the coupon MPCs are similar according to the bunching estimator (i.e., $MPC^{coupon}_{\tau} \approx MPC^{coupon}_{\tau'}$), then the formula above implies that the MPC estimated by comparing pairs of randomly assigned coupons should be approximately equal to the individual coupon MPCs.

There is a useful economic implication of the identity above, which is that if two coupons are equally cost-effective based on the coupon MPC estimate, then this implies that the government would have been able to increase total spending by randomly assigning a greater share of the coupons to the coupon with the higher threshold.

5 Main Results

We begin by presenting visual evidence of sharp bunching at the coupon-specific thresholds during the coupon wave but not during other time periods. We then turn to quantifying the effects of the coupons on spending. We estimate effects of coupons on targeted spending categories, non-targeted spending categories, and overall spending, and we also assess the importance of intertemporal substitution by calculating how the coupon MPCs evolve dynamically over time.

5.1 Graphical Evidence

We begin by presenting clear visual evidence of bunching at the coupon-specific thresholds. Recall that our analysis sample is the set of residents who acquired the coupons, tracking all of their spending on the platform in the periods before, during, and after the coupons were distributed.

For concreteness, we focus on the 54-18 life service coupon offered to City A residents in the second coupon wave. Panel A of Figure 2 shows the transaction-level spending distribution for recipients of this coupon in the two periods before the coupons were distributed. The similarity between the two pre-period distributions shows that there are no potentially confounding trends in overall spending happening in the pre-periods.

Next, Panel B of Figure 2 shows the spending distribution in the coupon wave period (t) relative to the pre-period ($t-1$). This figure shows clear visual evidence of "bunching" at the coupon specific threshold, which is consistent with coupon redemptions leading to an increase in spending. Moreover, to the left of the coupon-specific threshold, there is some visual evidence of "missing mass", which implies that some consumers are spending more than they otherwise would have in order to be able to redeem the coupon and earn the discount.

Lastly, Panel C of Figure 2 compares the spending distributions for the period following the coupon wave ($t+1$) to the pre-period ($t-1$), and the distributions are fairly similar, with perhaps some very slight evidence of fewer transactions across the distribution, which would be consistent with a very small amount of intertemporal substitution.

All of these results preview the main bunching estimates. As described in the previous section, these results can be used to reject the null hypothesis of no behavioral response to the coupons. Using the pre-period spending distribution as a counterfactual, we conclude that consumers using the coupons appear to be spending more than they otherwise would have. Figure 3 shows that the pre-period distributions are quite stable for several periods in a row leading up to the coupon wave, which means that our results are not sensitive to which pre-period we choose or whether we take an average of several pre-periods.

Figure 4 shows separate estimates of the "excess mass" and "missing mass" estimates for the 54-18 coupons by taking the difference between the period t and $t-1$ spending distributions. This figure also provides an additional way of assessing the validity of the research design, by assessing whether there is any positive (or negative) mass in the upper right tail of the spending distribution. We label this area the "excluded region" because it does not contribute to either the excess mass or the missing mass estimates. If there were confounding trends that affected the entire spending distribution, then this would lead to a positive or negative area in the excluded region. Since there is no mass in the excluded region, this gives us more confidence that the difference in distributions reflects the genuine causal effect of the coupons.

The Online Appendix reports the analogous figures for the other coupons in the data, and the same patterns tend to emerge: clear visual evidence of bunching at the coupon thresholds, excess mass that is much larger than the missing mass, and no mass in the excluded region in the upper right tail (Appendix Figures OA.1 - OA.15)

5.2 Empirical Estimates of MPC^{coupon}

To quantify the spending effects of the coupons, we estimate equation 1 to recover the coupon MPCs for each coupon. The results are reported in Table 3, which shows that the estimated MPCs range from 1.96 to 4.14, with a weighted average of 3.49-3.61.

As discussed above, these MPC estimates are fairly similar to the two existing estimates of coupon MPCs in the literature, and these estimates are also much larger than the estimates of other stimulus programs. Figure 5 compares the coupon MPCs from China's digital coupon program to MPCs from other programs such as tax rebates, stimulus payments, and other shopping coupon programs that did not feature spending thresholds like the China coupons. The main takeaway from Figure 5 is that the coupon MPCs are unusually large compared to other forms of fiscal stimulus.

To try to understand the reasons for the large estimated coupon MPCs, we first consider two possible explanations: (1) substitution between spending in the targeted category (i.e., the spending category targeted by the specific coupon) and spending in other categories; and (2) intertemporal substitution.

5.2.1 Substitution Between Spending Categories

Since we observe all of the spending on the platform for all of the consumers in our analysis sample, we can also calculate the coupon MPCs for spending in the non-targeted spending categories. For example, if we are looking at a supermarket coupon, we can look for evidence of substitution away from spending on food delivery or restaurants, or entertainment spending. Table 4 shows that we find no evidence of any statistically or economically significant effects of coupons on the spending in non-targeted spending categories.

This implies that if we aggregate total spending across all spending categories and re-estimate equation 1 using all platform spending rather than spending in the targeted category, we should find similar results. Figure 6 shows graphically that we see the same pattern of results for all spending for the same 54-18 coupon, and Table 5 shows that we find similar coupon MPCs when looking at total platform spending. This rules out substitution to other spending categories on the platform as an explanation for the large coupon MPCs.

5.2.2 Intertemporal Substitution

We next assess the role of short-run intertemporal substitution by re-estimating equation 1 for multiple additional periods before and after the coupons are distributed. These results are shown in Figure 7, which shows no evidence of substantial intertemporal substitution. There is a very slight decrease in spending in the following period, but this is nowhere near large enough to offset the substantial increase in spending in the coupon wave.

As a result, we conclude that the large coupon MPCs do not come primarily for substitution across different spending categories or from short-run intertemporal substitution. We thus conclude that the coupons successfully increased total consumer spending for at least a couple of months, as consumers increased spending to reach the coupon thresholds.

5.2.3 Robustness and Heterogeneity

We assess the robustness of these main results in two main ways. First, we re-estimate equation 1 with different values of H (which is the "tuning parameter" for the bunching estimator). We find very similar results for different values of H, which is not very surprising given the visual evidence presented above.We also estimate coupon MPCs exploiting the strict random assignment of coupons in the subset of coupons in our data. Figure 8 Panel A shows the extremely similar pre-period spending distributions for consumers who were randomly assigned different coupons. The similarity in distributions is consistent with strict random assignment. Panel B of Figure 8 shows that during the coupon wave the sharp bunching lines up with the coupon threshold for each subsample. Using the identity in equation 3 above, we find that the estimates based on strict random assignment are always very close to the implied estimates from the bunching estimates. The differences in magnitudes are always less than 10 percent, which supports the validity of the bunching estimates.

We also explore heterogeneity across consumers. We divided consumers into two roughly equal groups by age (above and below 35), and we find similar MPC estimates across the two groups (see Appendix Figure OA.20). We also divide consumers based on how often they used the platform prior to the coupon waves, and show that our baseline estimates are similar to estimates in the population of users who were active on the platform during our pre-periods. Somewhat mechanically, estimates of the MPC for users who were not active on the platform would be even higher than our baseline estimates (Appendix Figure OA.21). The fact that we also find broadly similar coupon MPCs for the "frequent users" (who are defined as consumers in the top 5 percent of spending distribution in the months prior to coupon issuance) partially addresses concerns about possible bias from not measuring "offline" consumption if this subsample contains a greater share of consumers with a large share of their overall spending in the targeted spending category observable on the platform.

Overall, our results consistently point toward large estimated coupon MPCs that do not come primarily from reduced spending in other categories or from short-run intertemporal substitution. Why then are the estimated coupon MPCs so large? The next section goes through a simple theoretical model of consumer spending to understand the economics behind these results.

6 Consumption Model Calibration

In this section, we begin by presenting a simple graphical model to reassess the simple economics of notches versus taxes. We view these graphs as part of our theoretical contribution, since they clarify and sharpen some of the intuition developed in previous work on notches.

We then lay out a simple dynamic consumption model to understand the economic determinants of the coupon MPC, and we compare it to the MPC out of an equivalent amount of cash provided by the government. We also use the model to compare coupons to a (linear) tax subsidy.

Lastly, we calibrate our consumption model and show that it can reproduce our empirical results, and we simulate alternative coupon designs using the calibrated model, compare coupons to taxes, and highlight when a social planner that cares about both increasing consumer welfare and increasing consumption in a particular sector will choose cash, coupons, linear tax subsidies, or some combination of each of these instruments.

6.1 Reassessing the Simple Economics of Notches Versus Taxes

In an earlier paper on tax notches, Blinder and Rosen (1985) describe the government as trying to stimulate consumption of a given commodity (e.g., by subsidizing charitable contributions through a linear tax subsidy). We adopt their single representative agent framework in this subsection to reassess some of the simple economics of notches compared to linear subsidies.

Panel A of Figure 9 shows a consumer allocating spending between spending on good A and good B. Before introducing the linear subsidy, the consumer chooses c_A^* and c_B^*. When the government introduces a linear subsidy (τ) on good A reducing the price from p to $p(1-\tau)$, this rotates out the consumer's budget constraint, which leads to higher consumer welfare and new choices $c_A^{'}$ and $c_B^{'}$. The total cost to the government from this subsidy is given by the vertical distance ON; to see this, note that the consumer would have spent D on good B had they chosen $c_A^{'}$ given their original budget constraint, and since the only change to budget constraint is the subsidy, this means the vertical distance is the cost of the subsidy which is $\tau * c_A^i = ON$.

Blinder and Rosen (1985) point out that the government could have instead designed a notch-based incentive where the notch is set at $c_A^{'}$ and transfers an amount ON in cash if the consumer chooses a level of consumption in sector A at or above the notch. This is shown in Panel B of Figure 9. Blinder and Rosen then note that the notch has created the same outcome for the consumer at the same revenue cost to the government. They write:

> *"The notch and linear schemes have the same revenue cost and induce the same behavior ... This example illustrates an obvious point. As long as one individual is being considered ... then there is nothing to choose between a linear incentive and a notch incentive."* (Blinder and Rosen 1985, p. 737)

We now show in this same graphical model that this reasoning is neither obvious nor accurate. The simple explanation is that while the reasoning in Blinder-Rosen is correct that a notch can always be designed to exactly replicate a linear subsidy incentive, the converse need not necessarily hold. In particular, the government can design a notch incentive that cannot be exactly replicated by any linear subsidy, because achieving the same increase in consumption in sector A would not have the same revenue cost and would not have the same effect on consumer welfare.

To see this graphically, Panel C of Figure 9 shows the government holding constant the cash transfer (at amount ON) but increasing the value of the notch. Panel C shows that the government can continue to increase the location of the notch up to the point where the consumer increases consumption in sector A to $c_A^{''}$ and the consumer is indifferent between increasing consumption to the notch (and receiving cash ON) and ignoring the notch.

Finally, Panel D of Figure 9 shows the linear subsidy the government would need to choose to achieve the same increase in consumption (from c_A^* to $c_A^{''}$). This subsidy is both more costly to the government (with vertical distance larger than ON), and the consumer strictly prefers the postsubsidy to the initial endowment, while the notch was designed to increase consumption in sector A while leaving the consumer exactly indifferent.

These simple figures illustrate the general point that while the government can design a notch to replicate a given linear subsidy and achieve the same outcome for the consumer at the same revenue cost, the government *cannot* replicate every notch with a linear subsidy. That is perhaps also an "obvious" point: the government has two parameters for the notch (the location of the notch and the vertical jump at the threshold) and only one parameter in the case of a linear subsidy, so it is not surprising that the two parameters can be used to replicate any linear subsidy but there can be two-parameter notches that cannot be exactly replicated by a linear subsidy.[10]

Overall, this graphical analysis shows that the government can create a notch that is more cost-effective (in terms of achieving a given increase in consumption per dollar of government revenue) but yields a smaller (or even no) change in consumer welfare – even in the simplified single-agent model studied in Blinder and Rosen. This highlights the key trade-off for policy: depending on how much the government cares about increasing consumer welfare relative to the policy-induced increase in consumption in the targeted sector, the government may strictly prefer using a notch rather than a linear subsidy.

We now use these results to calibrate a dynamic consumption model to try to quantify the trade-offs between cash, coupons, and linear subsidies.

6.2 Consumption Model Setup

The model is a T-period model with perfect foresight, no uncertainty, and exogenous income. Consumers can freely borrow, save, and allocate consumption across time periods and sectors (c_t^A and c_t^B). The consumer's per-period utility function is given by the following:

$$u(c_t^A, c_t^B) \equiv \frac{1}{1-\gamma} \left(\alpha(c_t^A)^\rho + (1-\alpha)(c_t^B)^\rho\right)^{(1-\gamma)/\rho}$$

where $\sigma = 1/(1-\rho)$ is the consumer's elasticity of substitution between consumption in sectors A and B, $1/\gamma$ is the intertemporal elasticity of substitution, and α is a share parameter that (along with σ) determines the share of spending in each period devoted to sector A.

[10] In fact, if we combine the linear subsidy with a lump-sum tax, then we can "fix" the claim in Blinder-Rosen because in the single-agent case there is "nothing to choose" between a linear incentive with a lump-sum tax and a notch incentive. This can be seen by vertically shifting down the τ'' line in Panel D of Figure 9 to intersect with the notch point.

The consumer's lifetime utility function is given by the following:

$$U \equiv u(c_1^A, c_1^B) + \frac{1}{1+\delta} u(c_2^A, c_2^B) + \ldots + \frac{1}{(1+\delta)^{T-1}} u(c_T^A, c_T^B)$$

The consumer maximizes lifetime utility subject to the following lifetime budget constraint:

$$c_1^A + c_1^B + \frac{c_2^A + c_2^B}{1+r} + \ldots + \frac{c_T^A + c_T^B}{(1+r)^{T-1}} \leq \sum_{t=1}^{T} \frac{y_t}{(1+r)^{t-1}}$$

where δ is the consumer's subjective discount rate, r is the exogenous interest rate, and $y_1, y_2, \ldots, y_T$ is the consumer's exogenous income.

6.3 MPC^{cash} versus MPC^{coupon}

If the government distributes cash in period 1 to the consumer, this is equivalent to an exogenous increase in y_1. In this case, we define $MPC^{cash} = \Delta(c_1^A + c_1^B)/\Delta(y_1)$. This can be solved in closed-form:

$$MPC^{cash} = \frac{1 - \frac{1}{1+r}\left(\frac{1+r}{1+\delta}\right)^{\frac{1}{\gamma}}}{1 - \left(\frac{1}{1+r}\left(\frac{1+r}{1+\delta}\right)^{\frac{1}{\gamma}}\right)^T}$$

Note that MPC^{cash} depends on r, δ, T, and γ, but does not depend on σ. Additionally, it is straightforward to show that if $r = 0 = \delta$, then $MPC^{cash} = 1/T$ as would be expected under perfect consumption smoothing.[11]

Now consider the government offering a coupon that offers ¥d off if the consumer spends more than ¥D in sector A. The consumer can decide whether or not to take up the coupon, but it must be used in period 1 (i.e., the coupon is "use it or lose it"). We assume the consumer takes up and uses the coupon if and only if it increases their lifetime utility. If the consumer takes up and uses the coupon, we can define $MPC^{coupon} = \Delta(c_1^A + c_1^B)/d$.

We cannot solve for $MPC^{coupon} = \Delta(c_1^A + c_1^B)/d$ in closed-form in general, but we prove in the Appendix that if $D < c_1^A + c_1^B$ then $MPC^{coupon} = MPC^{cash}$. In this case, the coupon is fungible with cash, and so the equivalence comes immediately from fungibility. If $D > c_1^A + c_1^B$, then the consumer may or may not take up the coupon, but if the consumer takes up the coupon then we

[11] To see this formally, take limits using L'Hôpital's rule as follows:

$$\lim_{r \to 0, \delta \to 0} MPC^{cash} = \lim_{x \to 1} \frac{1-x}{1-x^T} = \lim_{x \to 1} \frac{-1}{-Tx^{T-1}} = \frac{1}{T}.$$

expect $MPC^{coupon} > MPC^{cash}$.

If we define c_1^{A*} to be the optimal consumption choice in period 1 in sector A in the absence of the coupon, then if we assume that the consumers prefer to take up and use the coupon, then we can define the MPC^{coupon} as follows:

$$MPC^{coupon} = \frac{D - c_1^{A*}}{d} + \frac{\Delta(c_1^B)}{d}$$

This shows that if the change in consumption in sector B is small (i.e., $\Delta(c_1^B) \approx 0$), then $MPC^{coupon} \approx (D - c_1^{A*})/d$. This shows that the coupon MPC is increasing in the coupon threshold and decreasing in the coupon discount.

If the policymaker wants to maximize the "bang for back" of the coupon stimulus program, then this can be defined as maximizing MPC^{coupon} subject to the constraint that the consumer wants to take up and use the coupon. More formally, we can state this as choosing D to maximize MPC^{coupon} subject to constraint that $U^{coupon} > U^{baseline}$, and prove conditions under which this optimal value of D exists and prove in the Appendix that whenever the optimal value exists it is unique.[12]

We can also use the model to calculate the change in utility from receiving coupons compared to cash. Assuming that the coupon's threshold is large enough that it is not fungible with cash, we have the following approximation formula for the change in utility from receiving the coupon relative to the change in utility from receiving the equivalent amount in cash:

$$\frac{\Delta U^{coupon}}{\Delta U^{cash}} \approx 1 - 0.5 * (1 - \rho) \frac{(\Delta c_1^A)^2}{d * c_1^A}$$

The formula is derived by taking a second-order Taylor approximation around the utility from giving the consumer d in cash and then "forcing" the consumer to bunch at the coupon threshold. The envelope theorem allows us to ignore all other consumption changes, since those are re-optimized after the consumer bunches at the coupon threshold. The quadratic term comes from the second-order term, and it is scaled by $(1 - \rho)$; intuitively, if the consumer is very willing to substitute consumption between sectors, the consumer values the coupon almost as much as cash.

6.4 Model Calibration

We now show that we can reproduce the estimated coupon MPC quantitatively. We calibrate the model with the following parameters: $T = 10$, $y_t = 1$, $r = \delta = 0$, $\rho = 0.5$, $\gamma = 0.45$, and $\alpha = 0.25$.

[12] One easy way to see why D need not always exist is to recognize that if the threshold is greater than the consumer's lifetime income, there is no way to consumer to increase consumption in period 1 in sector A enough to use the coupon, so the coupon will not be used in this case.

The parameters imply that $c_t^A \approx 0.10\ \forall t$, and since there is no discounting and the interest rate is zero (i.e., $r = \delta = 0$), then we can see immediately that $MPC^{cash} = 1/T = 0.10$.

We solve for MPC^{coupon} numerically, and Figure 10 shows how the cash and coupon MPCs vary with the threshold (as the threshold goes from $D = 0.05$ to $D = 0.25$ and the coupon discount $d = 0.02$ is held constant throughout).

Specifically, Figure 10 shows that as long as the threshold is set below 0.10, the coupon MPC is equal to the cash MPC, which is a direct consequence of fungibility. As the threshold increases between 0.10 and 0.20, the coupon MPC increases approximately linearly. This figure also shows the change in consumer utility from taking up and using the coupon. The figure shows that eventually, the coupon's threshold is high enough that the consumer has higher utility from not using the coupon at all.

Intuitively, the difference between preferred consumption in sector A in period 1 is much lower than the consumption needed to achieve the coupon's discount, and this distorts consumer's preferred consumption bundle by enough that the positive effect of discount on consumer utility is not large enough to outweigh the utility loss from the distorted consumption. In the range of the coupon MPCs we estimate (1.5-3.7), the calibration above shows that this reduces consumer utility by about 50 percent relative to a benchmark where the government simply sent the coupon discount as cash (i.e., unconditional cash transfer).

Next, in Figure 11 (Panel A) we build on the same model simulation and decompose the coupon MPC into sector A and sector B spending. The figure shows that as the coupon MPC increases (as threshold increases), there is always very little substitution away from – or towards – sector B. In other words, the parameters chosen above lead to the coupon MPC being almost entirely driven by increased spending in sector A with the (net) increase in period 1 consumption coming almost entirely from this sector.

Lastly, in Panel B of Figure 11, we show that the previous result is sensitive to parameters. In particular, if we increase γ to 4 so that the intertemporal elasticity is 0.25, then we find that the coupon leads the consumer to reduce sector B consumption when using the coupon in period 1. Intuitively, this is because the consumer would rather substitute across spending categories than substitute consumption in both sectors over time. The coupon MPC ends up smaller because the total increased spending in sector A is a very large over-estimate of the total increase spending in period 1 because spending in sector B goes down substantially in the period the coupon is used. Our large coupon MPC estimates and lack of substitution away from other spending categories is consistent with a larger intertemporal elasticity.

6.5 Comparing Coupons to Linear Subsidies

We now compare the effects of a coupon to the effects of a temporary tax subsidy. In particular, we compare the effects of the coupon in the baseline calibration in the previous section to a temporary tax subsidy which introduces a subsidy τ_A in period 1 but not in any other periods. Following Blinder and Rosen (1985), we restrict ourselves to a linear subsidy and compare our coupon (with a notch) to this alternative policy instrument.

The graphical illustrations in Figure 9 show that the temporary tax subsidy can be replicated by a coupon in some cases, but that the converse need not hold (i.e., it may not be possible to replicate notch with subsidy). As a result, we expand the simulations in the previous section to introduce a linear subsidy, and we vary the tax from $\tau_A = 0$ to $\tau_A = 0.25$ in period 1. As we vary the tax, the consumer substitutes consumption towards sector A in period 1, and this comes from a combination of reduced spending in sector B in period 1 and reduced spending in all other periods (and the previous results show that for a modest intertemporal elasticity of substitution in consumption, we expect most of the increased in spending in sector A to come from future periods).

With the constant elasticity of substitution across categories and the constant intertemporal elasticity of substitution in consumption, it is straightforward to show that the linear subsidy MPC is approximately equal to $1/(1-\rho)$. In other words, for a given dollar of government spending on the linear subsidy, the planner can expect to get roughly $1/(1-\rho)$ additional dollars of spending in sector A. Figure 12 superimposes the effects of a linear tax subsidy on top of the effects of different coupons (varying the coupon thresholds as in Figure 11). Each circle shows the implied MPC and the effects of the linear subsidy on consumer welfare (relative to the effects of distributing the coupon as cash), and the circles move out to the right as the tax is increased from $\tau_A = 0$ to $\tau_A = 0.25$. When the tax is small, the welfare effects of the tax are similar to the effects of cash, which is a direct result of the envelope theorem, but the welfare effects are also small in absolute terms because the government is not spending very much when the tax rate is very low.

As the tax rate increases, eventually the government ends up spending the same amount on the tax as the government would have spent on the coupon simulated in Figures 10 and 11. We find this point is approximately $\tau_A \approx 0.147$. At this tax rate, the consumer is induced to spend just less than 0.14 on good A in period 1, which is approximately the same choice the consumer would make if the coupon with discount $d = 0.02$ was designed with a threshold at $D = 0.14$.

Once the tax rate is increased further, however, the MPC stays relative constant (falling only slightly), and consumer welfare increases because the government is continuing to increase the linear subsidy. This captures the government's trade-off: The stimulus is less cost-effective (based on MPC), but the increase in consumer welfare is higher.

We conclude by describing a stylized planner problem as follows: suppose the government can choose any mix of cash transfer, linear subsidy, and coupon (with coupon having arbitrary threshold and discount) subject to total government budget constraint B. Further, assume the government has a weight ω on the (money-metric) increase in consumer surplus and weight $(1-\omega)$ on the increase in spending in sector A caused by the stimulus policy (i.e., $\Delta(c_A)$). If $\omega = 1$, then the policymaker only cares about consumer surplus, and the optimal stimulus policy is to simply give a cash transfer equal to the total budget. If $\omega = 0$, then the policymaker only cares about increasing spending in sector A. In this case, we can show that the policymaker will never use the linear subsidy, but instead will choose a coupon that spends the entire budget on coupon discount and chooses the coupon threshold to maximize spending response ($\Delta(c_A)$) subject to incentive compatibility (i.e., that the consumer will still prefer to take up and use the coupon).

If $0 < \omega < 1$, then what we can say is that there exists a threshold value ω^*, where for all $\omega > \omega^*$ the policymaker will choose a mix of cash transfer and linear subsidy, or as in Blinder and Rosen (1985) the policymaker will replicate this outcome with a coupon that can achieve the same result.

Lastly, if $\omega < \omega^*$, then we can say that the policymaker will strictly prefer to use some of the total budget on a coupon (with spending threshold), because in this range it is not possible to replicate the policymaker's preferred coupon with a combination of a linear subsidy and a (non-negative) cash transfer.

To see the quantitative significance of these results, we simulate the same model in the previous subsection and search for all combinations of cash transfers, tax subsidies, and coupons to maximize social welfare according to this planner problem. We then can re-solve this planner problem for different levels of ω (at each value of $\omega \in \{0, 0.01, 0.02, \ldots, 0.98, 0.99, 1.0\}$). We find that for $\omega < 0.85$ the planner strictly prefers coupons to taxes. This implies that unless the planner puts a lot of weight on increasing consumer welfare through stimulus, any combination of cash and tax subsidies are dominated by a coupon.

These simulations show the potential for notch-based incentives to be preferred to cash transfers and linear tax subsidies. Obviously, the specific conclusions will be sensitive to calibrated parameters, and for simplicity we have worked completely in a single-agent/representative-agent setting. Blinder and Rosen (1985) already noted that with heterogeneity it is possible for notches to dominate linear subsidies depending on the parameterization of individual heterogeneity; by extensions that is likely to be the case in our planning problem, as well. What is novel in our setup is that even without heterogeneity we can see that there is a wide range of ω values where the planner strictly prefers coupons over tax subsidies, at least when the subsidies are restricted to linear subsidies and it is not possible to impose any lump-sum taxes.

7 Conclusion

This paper studies a novel form of targeted economic stimulus: government-issued digital coupons. These coupons were distributed across several provinces and municipalities in China in the aftermath of the Covid-19 recession, and in some cities, these coupons have become popular and continue to be distributed.

We use unique account-level and transaction-level data to estimate the effects of the coupons on consumer spending. We find that consumers increase spending by more than the discount they receive from the coupon. Our structured interviews with employees who worked for platform that distributed the coupons reveal that this was the intended goal of the coupons: to get consumers to spend more in the targeted spending categories, but have much of the increased spending come from the consumers themselves. That is why we call this *consumer-financed fiscal stimulus.*

Since the coupon thresholds were set fairly high, they create notches in the consumers' budget constraints above most consumers' usual spending amounts. The large MPCs we estimate can be interpreted as coming from large behavioral responses to notches. Since the coupon thresholds create notches, it is not surprising that coupons are less desirable than the equivalent amount distributed as cash. We estimate that consumers benefit from coupons, but consumers receive 50 percent of the increase in consumer welfare they would gain from cash.

If the policymakers are primarily interested in supporting targeted sectors, then it is easy to see from our model why coupons can have attractive targeting properties. Cash distributed by the government would mostly be spent on non-targeted sectors and saved for the future. The time-limited coupons, however, direct consumers to increase spending in the targeted sectors in order to receive the discount. Our model and calibrations capture the trade-off highlighted by the economist Justin Lin at the beginning of the paper: the more cost-effective the coupons are, the less valuable they are to consumers.

References

Berger, David, Nicholas Turner, and Eric Zwick, "Stimulating housing markets," *The Journal of Finance*, 2020, *75* (1), 277–321.

Best, Michael, James Cloyne, Ethan Ilzetzki, and Henrik Kleven, "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches," *Review of Economic Studies*, 2020, *87*, 656–690.

Blinder, Alan S. and Harvey S. Rosen, "Notches," *The American Economic Review*, 1985, *75* (4), 736–747.

Blomquist, Sören, Whitney K. Newey, Anil Kumar, and Che-Yuan Liang, "On Bunching and Identification of the Taxable Income Elasticity," *Journal of Political Economy*, 2021, *129* (8).

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-wage Jobs," *Quarterly Journal of Economics*, 2019.

Chetty, Raj, John N. Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team, "How did Covid-19 and Stabilization Polices Affect Spending and Employment?," *NBER Working Paper No. 27431*, 2020.

_ , _ , Tore Olsen, and Luigi Pistaferri, "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *Quarterly Journal of Economics*, 2011, *126* (2), 749–804.

Defusco, Anthony, Stephanie Johnson, and John Mondragon, "Regulating Household Leverage," *Review of Economic Studies*, 2020, *87.*

Government of the United Kingdom, "Eat Out to Help Out launches today-with government paying half on restaurant bills," www.gov.uk/government/news/eat-out-to-help-out-launches-today-with-government-paying-half-on-restaurant-bills 2020. Accessed: 2022-12-07.

Hsieh, Chang-Tai, Satoshi Shimizutani, and Masahiro Hori, "Did Japan's shopping coupon program increase spending?," *Journal of Public Economics*, 2010, *94*, 523–529.

Italian Ministry of Economy and Finance, "Italia Cashless," www.cashlessitalia.it. 2021. Accessed: 2022-12-07.

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles, "Household expenditure and the income tax rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

Kan, Kamhon, Shin-Kun Peng, and Ping Wang, "Understanding consumption behavior: evidence from consumers' reaction to shopping vouchers," *American Economic Journal: Economic Policy*, 2017, *9* (1), 137–53.

Kleven, Henrik, "Bunching," *Annual Review of Economics*, 2016, *8*, 435–464.

— **and Mazhar Waseem**, "Using Notches to Uncover Optimization Frictions and Structural Elasticites: Theory and Evidence from Pakistan," *Quarterly Journal of Economics*, 2013.

Liu, Qiao, Qiaowei Shen, Zhenghua Li, and Shu Chen, "Stimulating Consumption at Low Budget: Evidence from a Large-Scale Policy Experiment Amid the COVID-19 Pandemic," *Management Science*, 2021.

Mian, Atif and Amir Sufi, "The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program," *The Quarterly journal of economics*, 2012, *127* (3), 1107–1142.

Parker, Jonathan A, Jake Schild, Laura Erhard, and David Johnson, "Household spending responses to the economic impact payments of 2020: Evidence from the consumer expenditure survey," *NBER Working Paper No. 29648*, 2022.

Saez, Emmanuel, "Do Taxpayers Bunch at Kink Points?," *American Economic Journal: Economic Policy*, 2010, *2* (3).

Shapiro, Matthew D and Joel Slemrod, "Consumer Response to Tax Rebates," *American Economic Review*, 2003, *91* (1), 381–396.

— **and** — , "Did the 2008 tax rebates stimulate spending?," *American Economic Review*, 2009, *99* (2), 374–79.

Xing, Jianwei, Eric Zou, Zhentao Yin, Yong Wang, and Zhenhua Li, "'Quick Response' Economic Stimulus: The Effect of Small-Value Digital Coupons on Spending," *NBER Working Paper No. 27596*, 2021.

Table 1: Coupon Summary Statistics

City	Spending Category	Coupon Wave	Coupon	Coupons Available	Coupons Taken Up	Coupons Redemptions
City A	Supermarket	1	84-28	34,615	34,615	6,360
City A	Supermarket	1	54-18	46,154	46,154	7,552
City A	Supermarket	1	24-8	150,000	150,000	12,533
City A	Supermarket	2	84-28	130,150	73,243	19,850
City A	Supermarket	2	54-18	72,306	40,397	10,268
City A	Supermarket	2	24-8	86,767	49,024	7,639
City A	Life Services	2	114-38	42,066	20,732	9,499
City A	Life Services	2	84-28	41,955	20,735	10,230
City A	Life Services	2	54-18	194,069	94,690	49,626
City B	Food Delivery	1	30-15	100,000	7,198	2,688
City B	Food Delivery	2	30-15	46,000	46,000	5,006
City C	Life Services	1	100-40	40,179	40,179	26,004
City C	Life Services	1	200-100	6,000	6,000	5,332
City C	Life Services	2	100-40	19,814	19,814	13,062
City C	Life Services	2	200-100	3,000	3,000	2,520

Notes: This table gives detailed summary information on the coupons analyzed in our paper. “City” is the city that the coupon was available in. “Spending Category” is the category of spending in which the coupon could be redeemed. “Coupon Wave” is a (city-specific) number that sequences each release of coupons. “Coupon” displays the threshold and discount of the coupon. For example, a “24-8” coupon gives its holder 8 RMB off if they spend at least ¥24. “Coupons Available” is the number of coupons made available on our platform for the given coupon in the given wave period. “Coupons Taken Up” is the number of the coupons that were claimed by users of the platform. “Coupons Redeemed” is the number of coupons that were redeemed.

Table 2: Coupon Take-Up and Redemption Rates

City	Spending Category	Coupon Wave	Coupon	Take Up Rate	Redemption Rate
City A	Supermarket	1	84-28	1	0.18
City A	Supermarket	1	54-18	1	0.16
City A	Supermarket	1	24-8	1	0.08
City A	Supermarket	2	84-28	0.56	0.27
City A	Supermarket	2	54-18	0.56	0.25
City A	Supermarket	2	24-8	0.57	0.16
City A	Life Services	2	114-38	0.49	0.46
City A	Life Services	2	84-28	0.49	0.49
City A	Life Services	2	54-18	0.49	0.52
City B	Food Delivery	1	30-15	0.07	0.37
City B	Food Delivery	2	30-15	1	0.11
City C	Life Services	1	100-40	1	0.65
City C	Life Services	1	200-100	1	0.89
City C	Life Services	2	100-40	1	0.66
City C	Life Services	2	200-100	1	0.84

Notes: This table gives detailed summary information on the coupons analyzed in our paper. "City" is the city that the coupon was available in. "Spending Category" is the category of spending in which the coupon could be redeemed. "Coupon Wave" is a (city-specific) number that sequences each release of coupons. "Coupon" displays the threshold and discount of the coupon. For example, a "24-8" coupon gives its holder ¥8 off if they spend at least ¥24. "Take-Up Rate" is the fraction of coupons made available on our platform that were claimed by users of the platform. "Redemption Rate" is the fraction of taken-up coupons that were redeemed.

Table 3
Bunching Estimates of the Effects of Coupons on Spending in Targeted Spending Categories

City	Spending Category	Coupon Wave	Coupon	MPC^{coupon}
(1)	(2)	(3)	(4)	(5)
Panel A Coupon-Specific Coupon MPC Estimates				
City A	Life Services	2	54-18	3.05
				(0.03)
City A	Life Services	2	84-28	2.82
				(0.05)
City A	Life Services	2	114-38	2.37
				(0.04)
City A	Supermarket	1	24-8	4.14
				(0.09)
City A	Supermarket	1	54-18	3.61
				(0.05)
City A	Supermarket	1	84-28	3.32
				(0.05)
City A	Supermarket	2	24-8	3.94
				(0.09)
City A	Supermarket	2	54-18	3.82
				(0.04)
City A	Supermarket	2	84-28	3.50
				(0.03)
City B	Food Delivery	1	30-15	2.56
				(0.10)
City B	Food Delivery	2	30-15	1.96
				(0.07)
City C	Food Delivery	1	100-40	3.46
				(0.03)
City C	Food Delivery	1	200-100	1.98
				(0.03)
City C	Food Delivery	2	100-40	3.34
				(0.04)
City C	Food Delivery	2	200-100	2.05
				(0.04)
Panel B: Weighted-Average Coupon MPCs				
Weight by Number of Coupons Distributed				3.49
Weight by Number of Coupons Taken Up				3.61
Weight by Number of Coupons Redeemed				3.56

Notes: This table presents coupon MPC estimates using the bunching estimator described in the main text. Column (1) reports the anonymized city the coupon was distributed in, and columns (2) through (4) describe additional details of the coupon. Column (5) reports the coupon MPC estimate and the standard error in parentheses. All standard errors are heteroskedasticity-robust standard errors.

Table 4
Effects of Coupons on Spending in Other Spending Categories

	MPC estimates for each coupon and each spending category		
Coupon:	City A, Wave 2, 24-8	City A, Wave 2, 54-18	City A, Wave 2, 84-28
Spending Category:	(1)	(2)	(3)
Supermarket Spending	3.94	3.82	3.50
	(0.09)	(0.04)	(0.03)
All Other Spending	0.66	0.28	0.12
	(0.20)	(0.06)	(0.03)
Food Delivery	-0.81	-0.14	-0.15
	(0.14)	(0.04)	(0.02)
Dining	0.13	0.04	0.05
	(0.11)	(0.03)	(0.01)
Entertainment	0.13	0.00	0.02
	(0.05)	(0.02)	(0.01)
Hotel	0.06	0.03	0.00
	(0.04)	(0.01)	(0.01)
Shopping	0.05	0.02	0.01
	(0.01)	0.00	0.00
Movies	1.23	0.37	0.21
	(0.05)	(0.02)	(0.01)
Beauty	-0.13	-0.02	-0.02
	(0.03)	(0.01)	0.00
Total Spending	4.59	4.10	3.62
	(0.22)	(0.07)	(0.04)

Notes: This table presents coupon MPC estimates using the bunching estimator described in the main text for each coupon and category of spending separately, focusing on the Wave 2 coupons distributed in City A. All standard errors are heteroskedasticity-robust standard errors.

Table 5
Effects of Coupons on Total Platform Spending

City	Spending Category	Coupon Wave	Coupon	MPC^{coupon}
(1)	(2)	(3)	(4)	(5)
Panel A Coupon-Specific Coupon MPC Estimates				
City A	Life Services	2	54-18	3.10
				(0.03)
City A	Life Services	2	84-28	2.89
				(0.05)
City A	Life Services	2	114-38	2.42
				(0.04)
City A	Supermarket	2	24-8	4.59
				(0.22)
City A	Supermarket	2	54-18	4.10
				(0.07)
City A	Supermarket	2	84-28	3.62
				(0.04)
City B	Food Delivery	1	30-15	2.65
				(0.12)
City B	Food Delivery	2	30-15	2.13
				(0.08)
City C	Food Delivery	1	100-40	1.58
				(0.03)
City C	Food Delivery	1	200-100	1.13
				(0.03)
City C	Food Delivery	2	100-40	1.34
				(0.05)
City C	Food Delivery	2	200-100	1.39
				(0.05)
Panel B: Weighted-Average Coupon MPCs				
Weight by Number of Coupons Distributed				3.33
Weight by Number of Coupons Taken Up				3.45
Weight by Number of Coupons Redeemed				3.40

Notes: This table presents coupon MPC estimates using the bunching estimator described in the main text. Column (1) reports the anonymized city the coupon was distributed in, and columns (2) through (4) describe additional details of the coupon. Column (5) reports the coupon MPC estimate and the standard error in parentheses. All standard errors are heteroskedasticity-robust standard errors.

Table 6
Coupon MPC Heterogeneity by Age

City	Spending Category	Coupon Wave	Coupon	MPC^{coupon} estimate		
				Full sample	Age $\geq$ 35	Age < 35
(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A Coupon-Specific Coupon MPC Estimates						
City A	Life Services	2	54-18	3.05	3.13	3.01
				(0.03)	(0.05)	(0.04)
City A	Life Services	2	84-28	2.82	2.88	2.78
				(0.05)	(0.08)	(0.06)
City A	Life Services	2	114-38	2.37	2.52	2.29
				(0.04)	(0.07)	(0.05)
City A	Supermarket	1	24-8	4.14	4.77	3.48
				(0.09)	(0.12)	(0.12)
City A	Supermarket	1	54-18	3.61	3.69	3.50
				(0.05)	(0.07)	(0.08)
City A	Supermarket	1	84-28	3.32	3.32	3.31
				(0.05)	(0.06)	(0.07)
City A	Supermarket	2	24-8	3.94	4.13	3.72
				(0.09)	(0.13)	(0.13)
City A	Supermarket	2	54-18	3.82	3.81	3.83
				(0.04)	(0.06)	(0.07)
City A	Supermarket	2	84-28	3.50	3.46	3.55
				(0.03)	(0.03)	(0.04)
City C	Food Delivery	1	100-40	1.66	1.66	1.66
				(0.03)	(0.05)	(0.04)
City C	Food Delivery	1	200-100	1.31	1.13	1.39
				(0.04)	(0.08)	(0.05)
City C	Food Delivery	2	100-40	1.14	1.23	1.11
				(0.03)	(0.06)	(0.04)
City C	Food Delivery	2	200-100	1.37	1.52	1.31
				(0.05)	(0.09)	(0.06)

Notes: This table presents coupon MPC estimates using the bunching estimator described in the main text. Column (1) reports the anonymized city the coupon was distributed in, and columns (2) through (4) describe additional details of the coupon. Column (5) reports the coupon MPC estimate and the standard error in parentheses. All standard errors are heteroskedasticity-robust standard errors.

Table 7
Heterogeneity by Pre-Period Platform Usage

City	Spending Category	Coupon Wave	Coupon	MPC^{coupon} estimate			
				Full sample	Inactive users	Active users	Frequent users
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			Panel A Coupon-Specific Coupon MPC Estimates				
City A	Life Services	2	54-18	3.05	4.14	2.88	1.95
				(0.03)	(0.07)	(0.03)	(0.10)
City A	Life Services	2	84-28	2.82	3.59	2.69	2.21
				(0.05)	(0.11)	(0.06)	(0.17)
City A	Life Services	2	114-38	2.37	3.09	2.25	1.89
				(0.04)	(0.09)	(0.05)	(0.14)
City A	Supermarket	1	24-8	4.14	4.86	2.72	1.49
				(0.09)	(0.09)	(0.17)	(0.30)
City A	Supermarket	1	54-18	3.61	3.78	3.09	2.56
				(0.05)	(0.06)	(0.12)	(0.20)
City A	Supermarket	1	84-28	3.32	3.43	2.96	2.75
				(0.05)	(0.05)	(0.10)	(0.18)
City A	Supermarket	2	24-8	3.94	4.54	2.59	1.56
				(0.09)	(0.10)	(0.21)	(0.35)
City A	Supermarket	2	54-18	3.82	3.95	3.40	2.79
				(0.04)	(0.05)	(0.10)	(0.18)
City A	Supermarket	2	84-28	3.50	3.59	3.17	2.72
				(0.03)	(0.03)	(0.06)	(0.11)
City C	Food Delivery	1	100-40	1.66	2.06	1.50	1.15
				(0.03)	(0.05)	(0.04)	(0.11)
City C	Food Delivery	1	200-100	1.31	1.59	1.22	0.89
				(0.04)	(0.08)	(0.05)	(0.14)
City C	Food Delivery	2	100-40	1.14	1.51	1.00	0.74
				(0.03)	(0.06)	(0.04)	(0.10)
City C	Food Delivery	2	200-100	1.37	1.73	1.24	0.92
				(0.05)	(0.09)	(0.06)	(0.15)

Notes: This table presents coupon MPC estimates using the bunching estimator described in the main text. Column (1) reports the anonymized city the coupon was distributed in, and columns (2) through (4) describe additional details of the coupon. Column (5) reports the coupon MPC estimate and the standard error in parentheses. All standard errors are heteroskedasticity-robust standard errors.

Figure 1: Heterogeneity in Coupon Design: Variation in Thresholds and Discounts

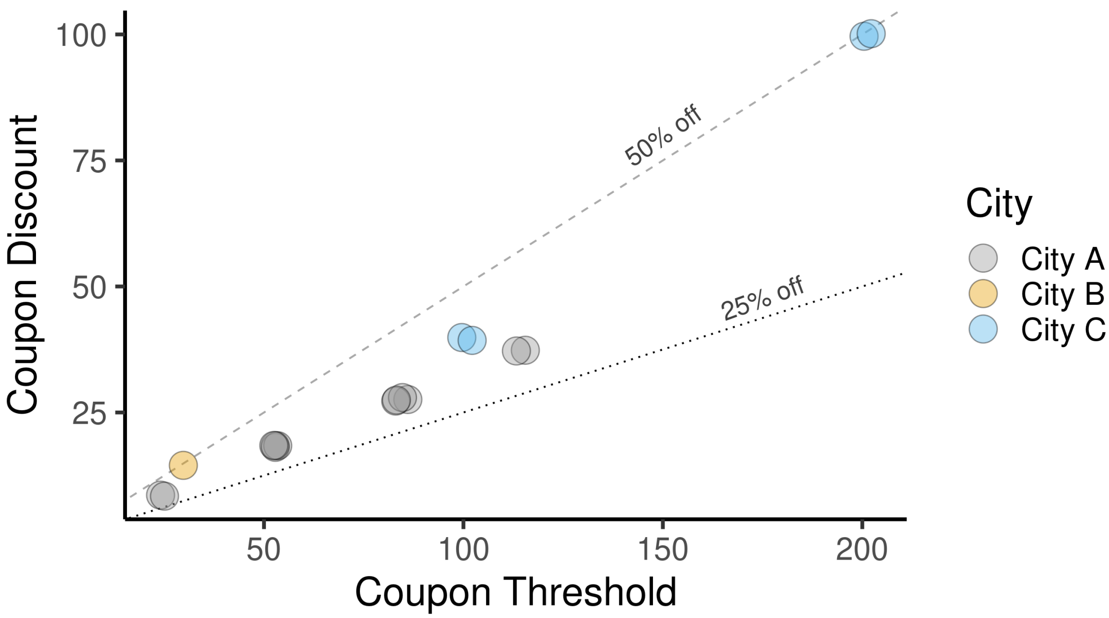

Notes: This figure shows the distribution of coupon thresholds and discounts in our data. The dashed lines indicate the set of coupon discounts that corresponds to 25 percent and 50 percent of the coupon thresholds. All of the coupons lie between the two rays, which implies that when municipalities chose higher coupons, they chose higher coupon discounts to keep the ratio of the discount to the threshold between 25 and 50 percent. All values are in ¥.

Figure 2: Illustration of Bunching Estimator for 54-18 Life Services Coupon in City A

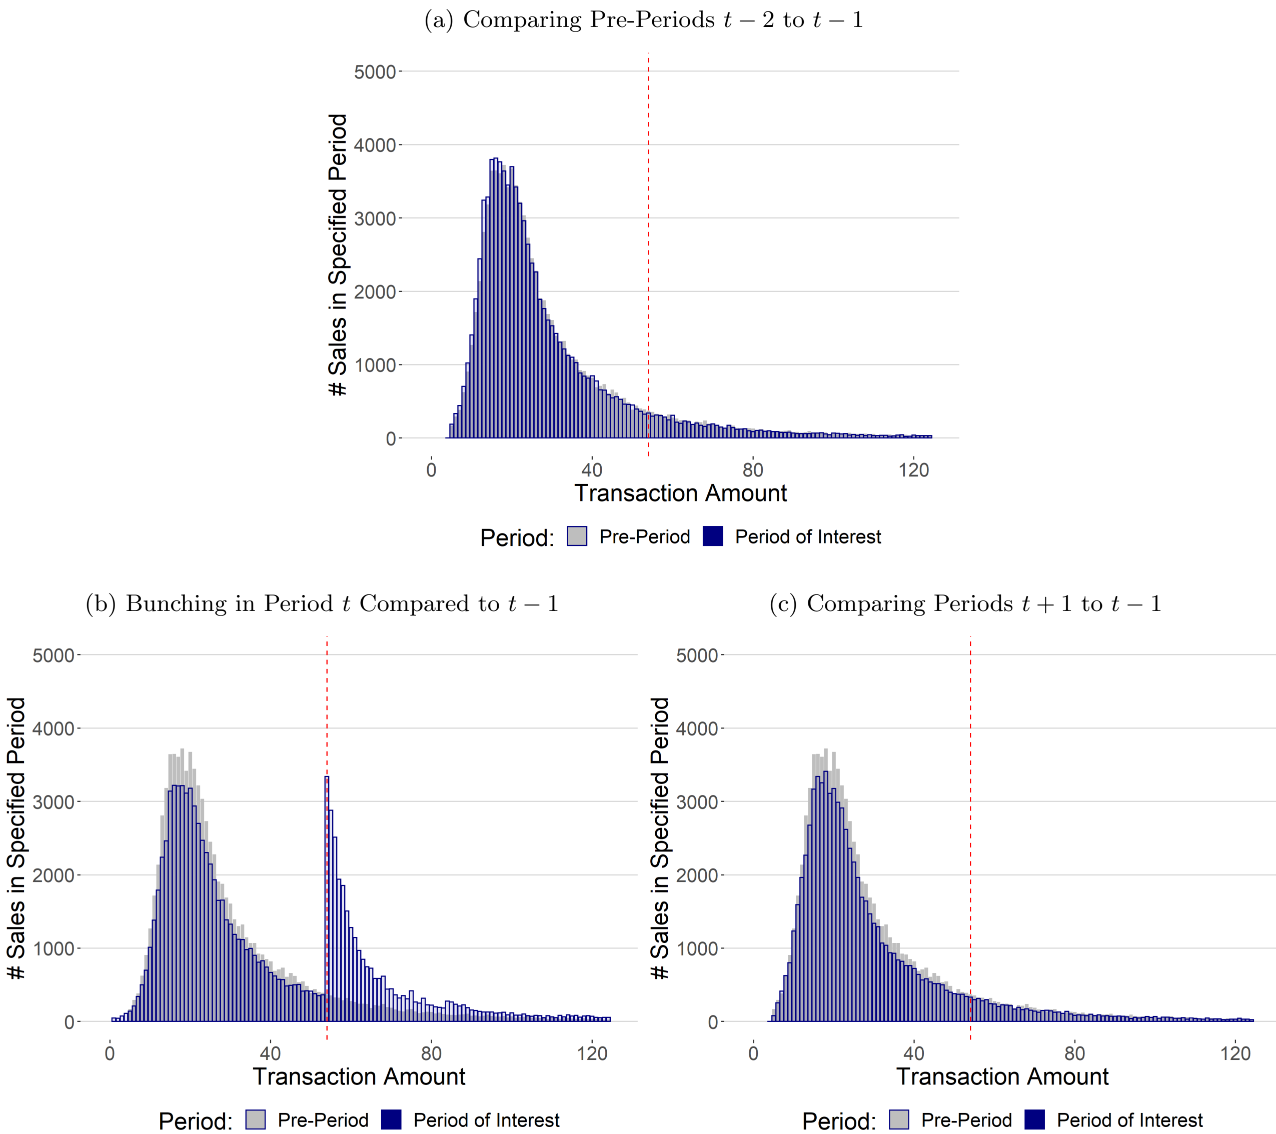

Notes: This figure illustrates the bunching estimator by comparing the distribution of spending between periods around the time the coupons were distributed. Panel (a) compares the distribution of spending in the two pre-periods immediately before the coupons were distributed. Panel (b) shows the distribution of spending during the coupon wave. Panel (c) shows the distribution of spending in the period immediately after coupons were distributed. In all panels the pre-period $t-1$ distribution is shown for reference.

Figure 3: Sensitivity of Bunching Estimator to Alternative Pre-Periods for 54-18 Coupon in City A

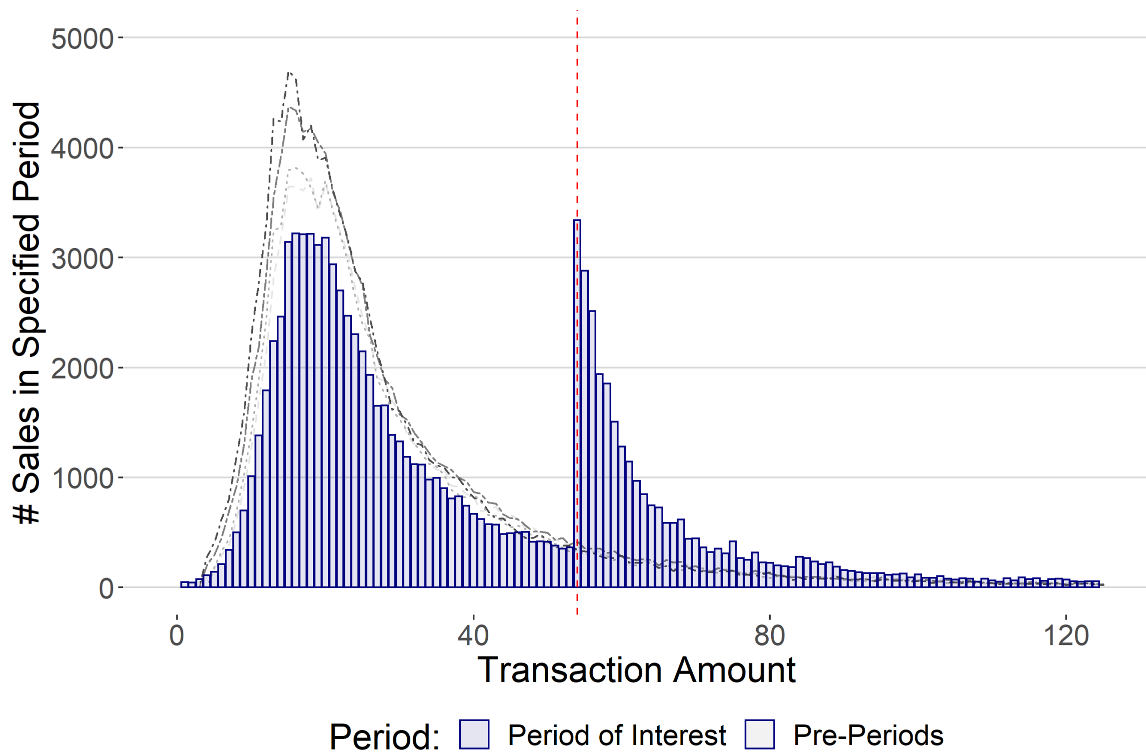

Notes: This figure illustrates the sensitivity of our bunching estimator to different pre-periods by comparing the distribution in the coupon wave period to seven pre-periods ($t-1$ through $t-7$).

Figure 4: Excess Mass and Missing Mass Estimates

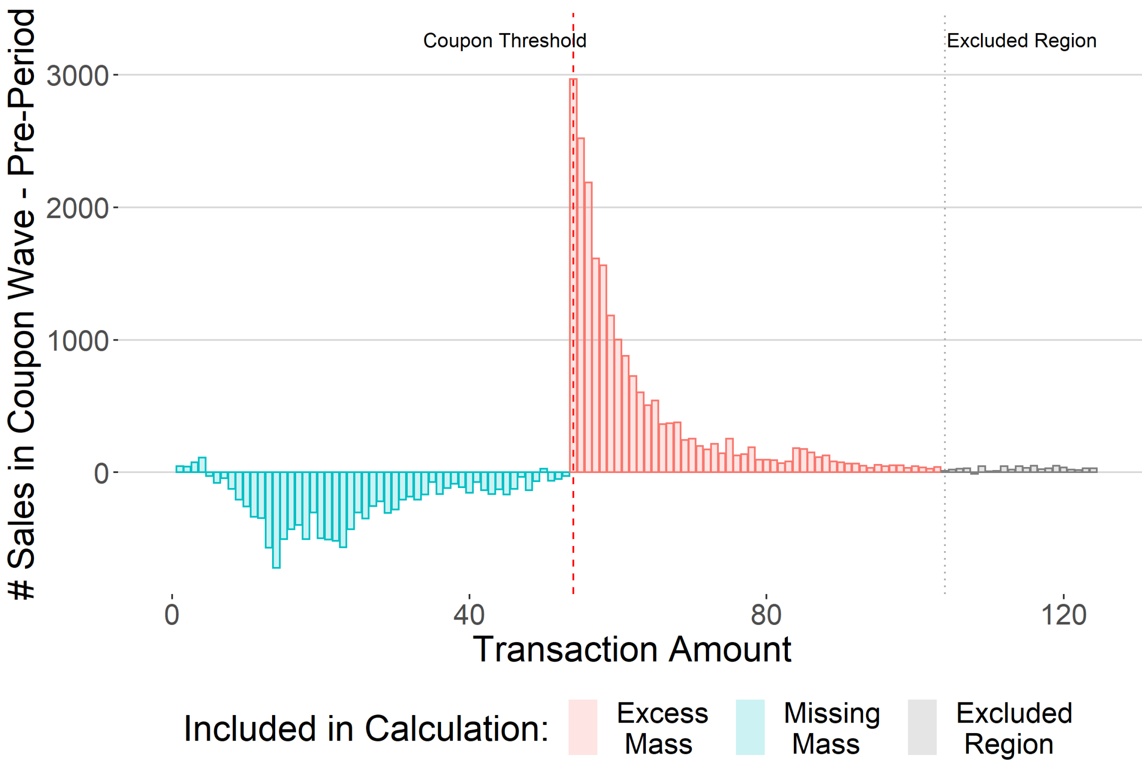

Notes: This figure "differences" the transaction distributions between the periods t and $t-1$ as shown in panel (b) of Figure 2. The area below the coupon threshold is the "missing mass" and the area above the threshold (τ) and below the assumed upper bound ($\tau + H$) is the "excess mass". The difference in distributions above the upper bound is the excluded region and is assumed to be equal to zero if the research design is valid.

Figure 5: Comparison with MPC Estimates from the Literature

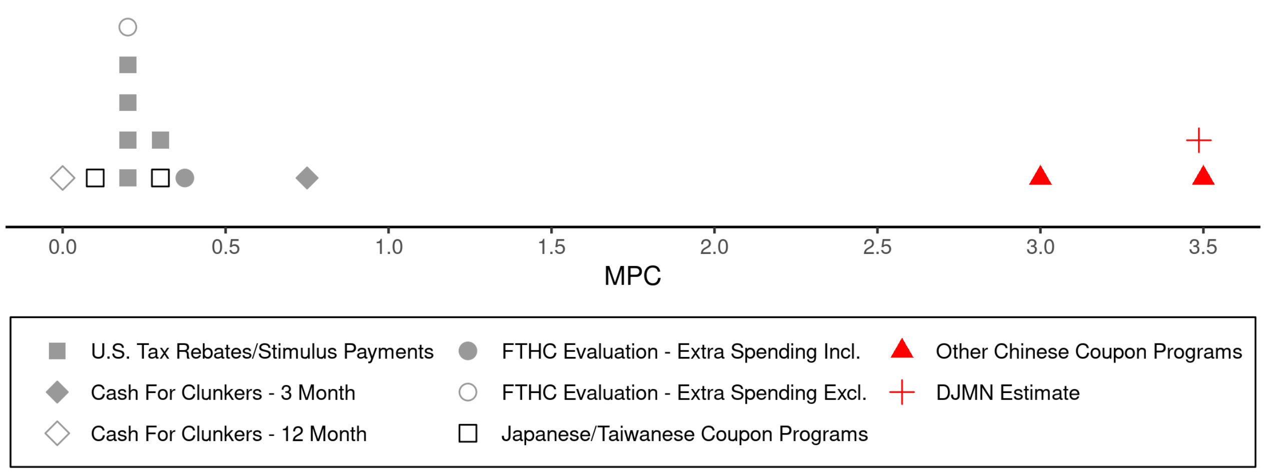

Notes: This figure contrasts our baseline MPC estimates to previous MPC estimates in the literature. The grey triangles represent MPC estimates from the US Cash for Clunkers program (Mian and Sufi 2012). The grey circles represent estimates from the US first-time homebuyer tax credits (Berger et al. 2020). The filled grey squares represent estimates from US tax rebates or cash stimulus payments (Shapiro and Slemrod 2003; Johnson et al. 2006; Shapiro and Slemrod 2009; Chetty et al. 2020). The empty squares with thick black borders represent estimates from coupon stimulus programs in Japan and Tawain which did not feature threshold designs (Hsieh et al. 2010; Kan et al. 2017). The red triangles represent estimates from other papers evaluating Chinese cities' Covid-era coupon programs (Liu et al. 2021; Xing et al. 2021). The red cross ("DJMN Estimate") represents our baseline estimates, aggregated across cities and waves.

Figure 6: Effects of Coupons on Total Platform Spending

(a) Bunching in Period t Compared to $t-1$ (b) Comparing Periods $t+1$ to $t-1$

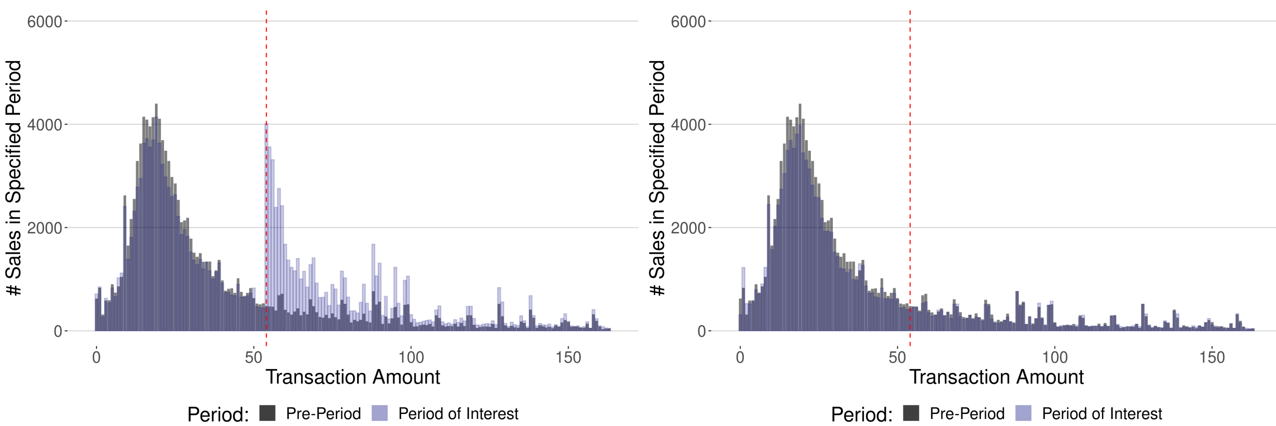

Notes: This figure reproduces the panels in Figure 2 using the distribution of total spending on the platform instead of the distribution of spending in the spending category targeted by the coupon. The similarity in figures across the analogous panels is consistent with the estimates in Table 5 showing limited effects of coupons on consumption in "non-targeted" spending categories.

Figure 7: Evolution of Coupon MPC estimates Over Time

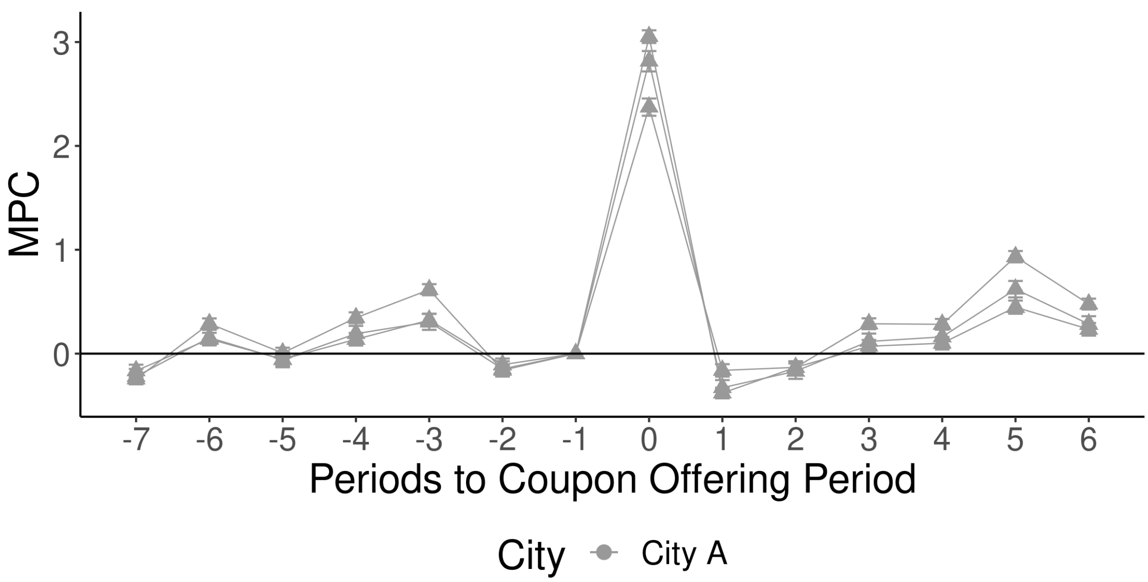

Notes: This figure reports coupon MPC estimates over time for the three “Life Service” coupons distributed in wave 2 in City A. The small coupon MPC estimates after the coupon wave period is consistent with a very small amount of short run intertemporal substitution.

Figure 8: Estimating MPC^{coupon} by Exploiting Random Assignment of Coupons

(a) Comparing Spending in Pre-Period $t-1$

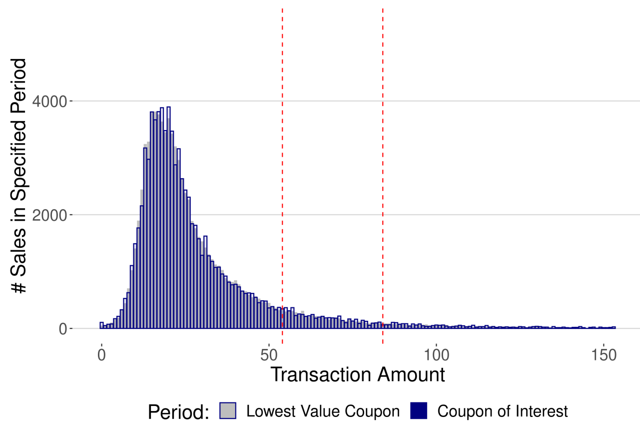

(b) Comparing Spending in Coupon Period t

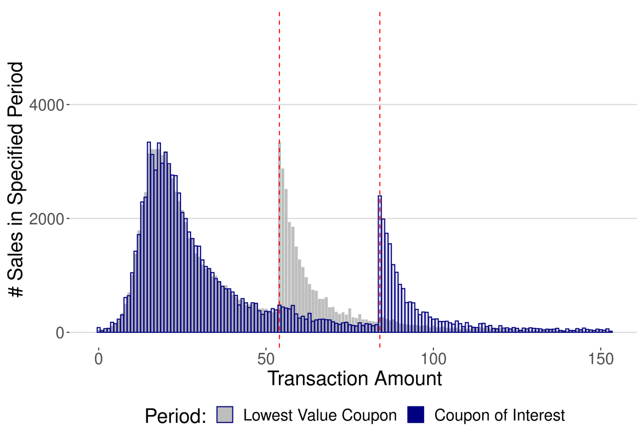

(c) Comparing Spending in Post-Period $t+1$

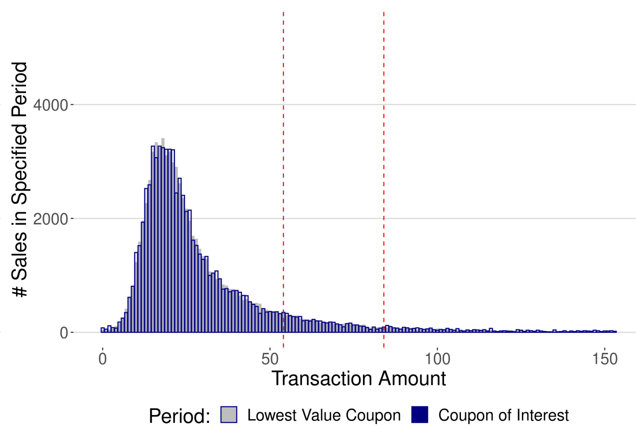

Notes: This figure reports panels analogous to Figure 2 except that the identification is based on comparing the consumers who were randomly assigned different coupons in Wave 2 in City A. Panel (a) compares the distribution of spending between the two groups of consumers assigned either the $54-18$ or the 84-28 coupon. The distributions are nearly identical which is consistent with the strict random assignment of the coupons. Panel (b) compares the distribution of spending during the coupon wave; there is clear bunching at the coupon thresholds for each group, and there is greater overall spending for the consumers randomly assigned the higher-threshold/higher-discount coupon. Panel (c) shows the distribution of spending in the period immediately after coupons were distributed; the similarity is consistent with limited amount of intertemporal substitution, since the greater spending in coupon wave period does not show up as lower spending in the following period.

Figure 9: Graphical Model

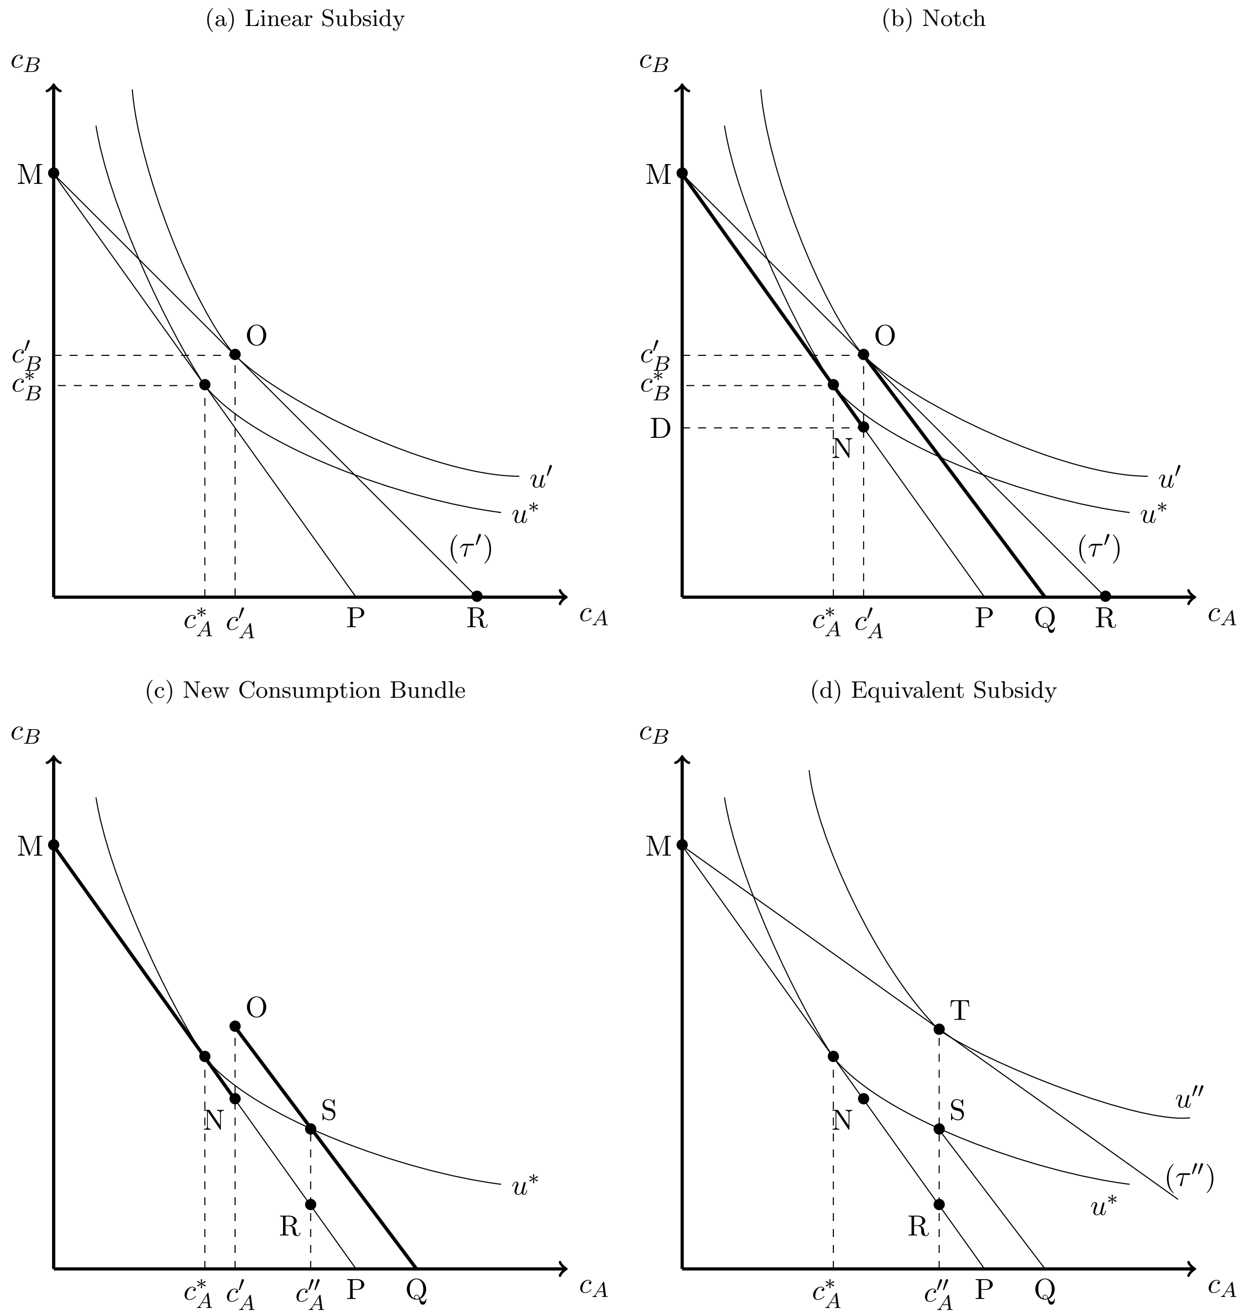

Notes: This figure presents a simple two-good graphical model to reassess the economics of notches versus linear subsidies. In Panel (a), the consumer responds to a linear subsidy that reduces the price of good A by a factor $(1 - \tau')$. This rotates the budget constraint and leads to new choices $c_A^{'}$ and $c_B^{'}$. Panel (b) shows that the government can replicate the outcome of the linear subsidy with a notch that transfers ON to the consumer if they choose at least c_A^i of good A. Panel (c) shows that the government can design a notch with a higher threshold where the consumer is indifferent between locating at the notch and remaining at initial endowment; this new notch has same cost to government ($ON = SR$), but leads to a large increase in consumption of good A. Lastly, Panel (d) shows the linear subsidy that is necessary to induce the consumer to increase consumption by same amount as in Panel (c). This shows that a linear subsidy is not equivalent to the notch, since to achieve the same increase in consumption of good A the linear subsidy leads to a greater increase in consumer welfare but also a larger amount of government spending (RT instead of RS).

Figure 10: Model Simulation of MPC^{coupon}

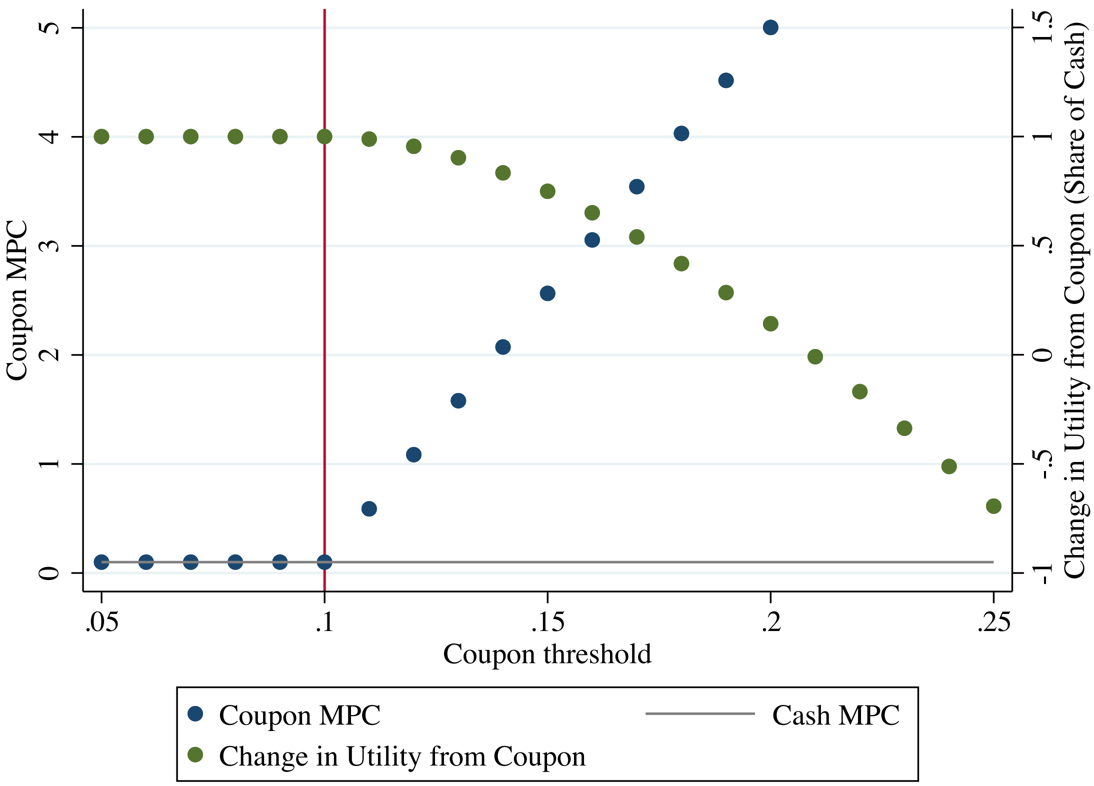

Notes: This figure shows how the model-based MPC^{coupon} varies with the coupon threshold. See main text for more details on the simulation.

Figure 11: Sensitivity in Model to Different Values of the Intertemporal Elasticity of Substitution

(a) Model Simulation with $\gamma = 0.5$

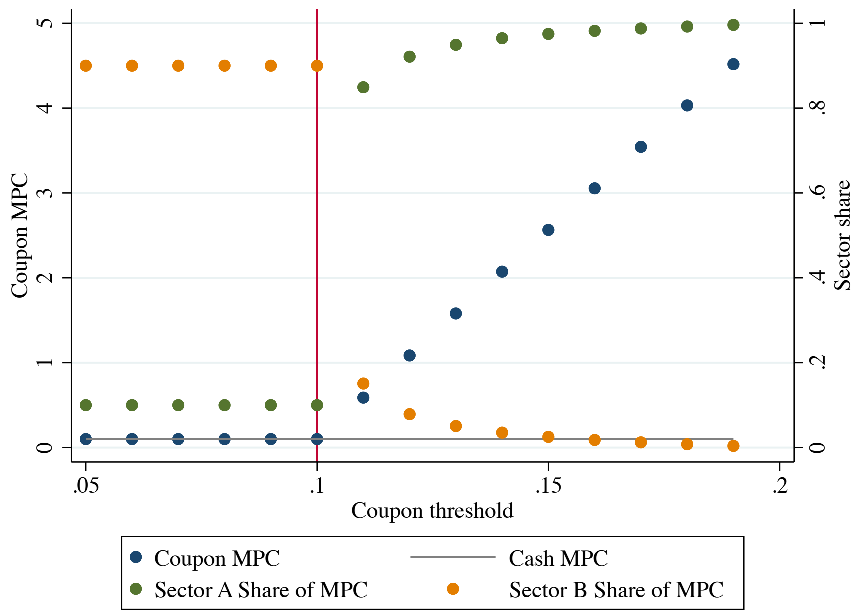

(b) Model Simulation with $\gamma = 0.5$

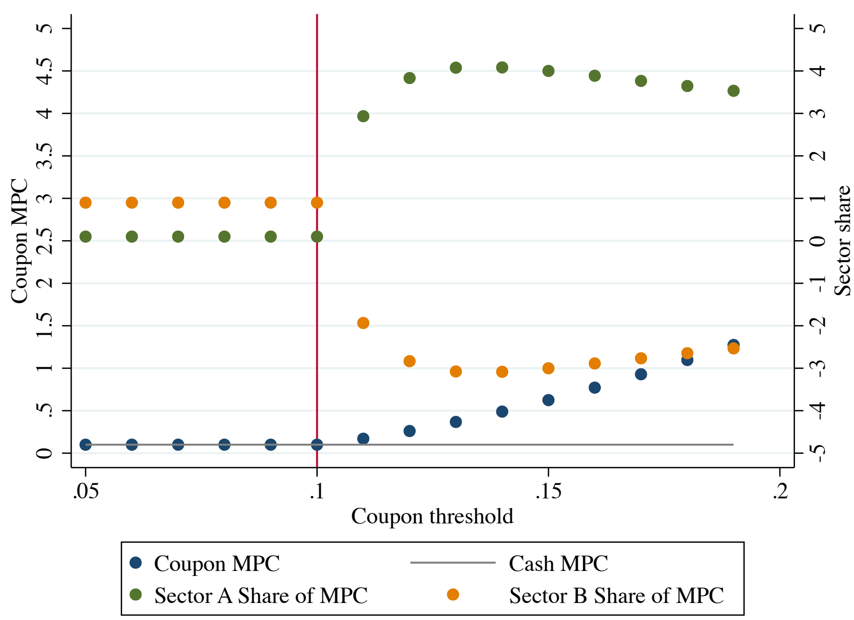

Notes: This figure shows how the model-based simulation results vary with the intertemporal elasticity of substitution $(1/\gamma)$.

Figure 12: Model Simulation for Varying Coupon Threshold and Tax Subsidy

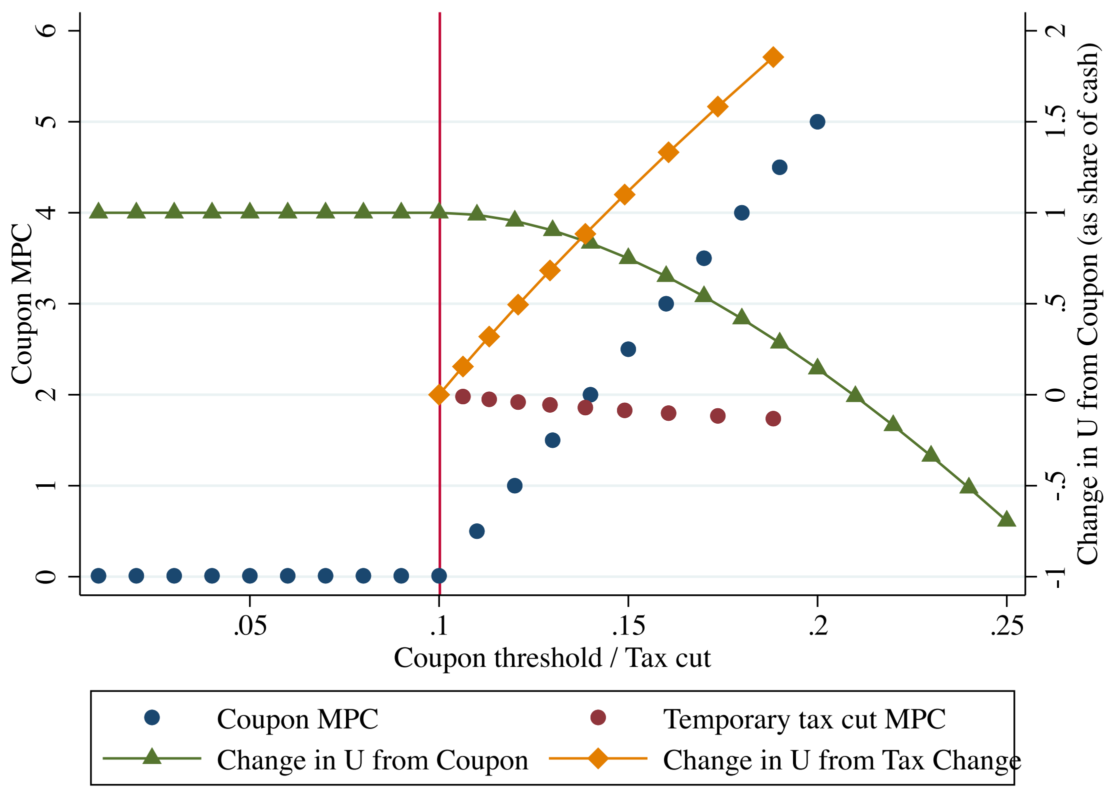

Notes: This figure reports the model-based simulation results from two policy scenarios: a coupon with a varying coupon threshold (and fixed discount), and a varying linear tax subsidy. The tax subsidy scenario considers values $\tau_A \in \{0.0, 0.03, 0.06, \ldots, 0.24, 0.27\}$. When $\tau_A \approx 0.15$, the tax subsidy and the coupon at threshold $D = 0.14$ achieve approximately the same outcome – specifically, the same effect on consumer utility and the same effect on the government budget (i.e., the same revenue cost).