

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

Jing Ding	Lei Jiang	Lucy Msall	Matthew J. Notowidigdo
Tongji University	Kent State University	University of Chicago	University of Chicago and NBER

First Version: February 25, 2022

This Version: April 17, 2024

Abstract

In 2020, local governments in China began issuing digital coupons to stimulate spending in targeted categories such as restaurants and supermarkets. Using data from a large e-commerce platform and a bunching estimation approach, we find that the coupons caused large increases in spending of 3.1–3.3 yuan per yuan spent by the government. The large spending responses do not come from substitution away from non-targeted spending categories or from short-run intertemporal substitution. To rationalize these results, we develop a dynamic consumption model showing how coupons' minimum spending thresholds create temporary notches that lead to large spending responses.

*We thank I-Han Lai, Angelo Marino, Jeremy Mopsick, and Ningxin Zhang for excellent research assistance and Milena Almagro, Emir Kamenica, Koichiro Ito, Zhiguo He, Nancy Qian, Peter Ganong, Anders Humlum, Kory Kroft, Maria Polyakova, 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 opinions, 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 or governments.

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

Many countries distribute stimulus payments during economic downturns to increase consumption. For example, the US government distributed stimulus payments to households in each of the last three recessions, and each time, households used the payments to immediately increase consumption (Johnson et al. 2006; Shapiro and Slemrod 2009; Parker et al. 2022). Many governments also design 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 housing market 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 stimulus using government-issued digital coupons. The coupons were delivered through smartphone apps and designed to encourage spending in certain sectors such as restaurants, grocery stores, and shopping malls. These sectors were hit particularly hard during the early months of the COVID-19 pandemic in China. The digital coupons had fixed spending thresholds that needed to be reached before consumers received money from the government—for example, a coupon would give 18 yuan off 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 from a large online platform covering several different types of coupons distributed across three cities in China. The data set includes the spending amount and the time and date of each transaction for everyone who received coupons. The different coupons have a range of different spending thresholds and apply to different spending categories. Throughout the paper, we define the “coupon MPC” (MPC^{coupon}) as the increase in consumption caused by a coupon relative to the coupon’s fiscal cost. For example, if 50,000 “Spend at least ¥54, get ¥18 off” food delivery coupons were used in a city, then the fiscal cost is $18 \times ¥50,000 = ¥900,000$. If the total increase in spending caused by the coupons is ¥1,800,000, then we would estimate $MPC^{coupon} = 2.0$.

As we describe in detail below, the reason that the MPC^{coupon} can be larger than one is that many consumers may need to increase their spending substantially in the targeted spending categories to

reach the spending threshold and take advantage of the coupon. In doing so, if they do not decrease their spending in other categories, then their total spending would increase by more than the discount associated with the coupon, which is the amount financed by the local government. Because of this, we call this new form of fiscal stimulus *consumer-financed fiscal stimulus* since whenever $MPC^{coupon} > 1$, the increased spending caused by the government is partly paid for by consumers.

We begin our empirical analysis by presenting clear visual evidence of sharp “bunching” at coupon-specific thresholds during the weeks that the coupons could be used. We find no evidence of similar bunching in the weeks before or after the coupons were distributed, indicating a clear behavioral response to the coupon-specific spending thresholds. We then use a bunching estimator following [Kleven \(2016\)](#) that compares the entire transaction-level spending distribution before and after the coupons were distributed. Under the assumption that the pre-period spending distribution is a valid counterfactual, we can identify and estimate the MPC^{coupon} coupon by coupon by integrating over the difference in spending distributions between periods.

Turning to our main results, we find a range of MPC^{coupon} estimates across coupons (1.9 to 4.6), with a weighted average of 3.1–3.3. We assess whether the large MPC^{coupon} estimates come from substitution between “targeted” and “non-targeted” spending categories using data on all consumer spending on the platform, and we find no evidence of meaningful cross-category substitution. We also find very little intertemporal substitution in the short run, with the MPC^{coupon} estimates remaining fairly stable for several months after the coupons were distributed. Our main results are robust to several alternative ways of estimating the “bunching” in the spending distribution, and we find similar results from an alternative empirical approach that exploits the explicit random assignment of coupons for a subset of the coupons in our data. As far as we know, this is the first time a bunching estimator is validated using explicit random assignment.

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 calibrate our model to match our MPC^{coupon} estimates and find that the key to matching our reduced-form results is that the coupon threshold must be set higher than the spending in the targeted sector that many consumers would have preferred in the absence of the coupon. This assumption appears to hold in our empirical setting, given the location of the threshold in the pre-period spending distribution.

We also use the calibrated model to illustrate how the MPC^{coupon} varies with the coupon’s threshold and calculate the welfare cost to consumers from receiving a coupon instead of cash. We find that consumers obtain approximately 50 percent of the increase in consumer welfare that they would have received from an equivalent amount of fiscal stimulus distributed as cash but that spending increases much more in the targeted sectors with coupons than with cash, highlighting the potentially attractive targeting properties of coupons as stimulus.

Taken together, our empirical and theoretical results suggest that digital coupons are a cost-

effective way to provide stimulus targeted to specific sectors.¹ The MPC^{coupon} estimates are large, and the effects persisted for several months, implying that the increased spending from the coupons is achieved at a very low fiscal cost relative to other forms of stimulus.

Our paper contributes to three main areas of research. First, we contribute to the study of consumption responses to fiscal stimulus. This literature includes the stimulus papers mentioned above and recent related work studying shopping coupons in Japan and shopping vouchers in Taiwan (Kan et al. 2017; Hsieh et al. 2010).

Second, our paper contributes to the study of tax notches, building on the early work by Blinder and Rosen (1985). We correct a small inaccuracy in their analysis of when linear incentives and notches are equivalent, and our correction shows that notches may be strictly preferable to linear subsidies in a broader range of settings than previously recognized. Our empirical approach is broadly related previous work that uses “bunching” to infer behavioral responses to tax kinks, tax notches, and minimum wages (Best et al. 2020; Defusco et al. 2020; Cengiz et al. 2019; Kleven and Waseem 2013).

Last, our paper is most closely related to two other recent studies of digital coupons in China using different data sets and empirical approaches. Both papers report estimates that are broadly similar to our main results despite different data and research designs. Xing et al. (2021) study digital coupons in a single large Chinese city and estimate an average MPC^{coupon} of approximately 3.0 by comparing “near-miss” consumers who just barely missed out on receiving a coupon to consumers who just barely received a coupon.² Liu et al. (2021) use administrative data on coupons issued on Alibaba in Hangzhou and Guangxi and use a difference-in-difference approach comparing consumers who received coupons to a random sample of individuals who tried but failed to obtain a coupon. They report MPC^{coupon} estimates in the range of 3.4–5.8. Relative to the analyses in these papers, ours covers a larger number of cities and coupons and a wider range of coupon thresholds and discounts. We also exploit the explicit random assignment of coupon thresholds and discounts, which is unique to our setting. Our bunching estimator approach can also be used for all the coupons in our data, while the “near-miss” research design in Xing et al. (2021) is infeasible to implement for the coupons in our data for which take-up was incomplete (which is the case for 7 of the 15 coupons in our data). Finally, unlike the previous two papers, we develop and calibrate a model that we use to compare the consumer welfare effects of coupons and cash, compare coupons with temporary tax subsidies, and evaluate counterfactual coupon designs.

¹Throughout our paper, we take as given the policymaker’s objective of increasing spending in the short run in particular sectors. Prior work in macroeconomics has identified situations when temporary tax changes can be useful (Correia et al. 2008, 2013), but the analyses have focused on state-specific rather than sector-specific tax instruments. We conjecture that the recent analysis of “Keynesian supply shocks” during a pandemic (see, e.g., Guerrieri et al. 2022) can be extended to provide a more rigorous justification for when a policymaker would want to provide a targeted temporary tax cut to a specific sector. If so, then our analysis suggests that in some settings it may be preferable for the policymaker to use temporary notches rather than temporary tax subsidies to increase spending in particular sectors.

²Xing et al. (2021) also estimate 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.

2 Background and Data

2.1 Background on the Chinese Coupon Programs

In response to the COVID-19 pandemic, which slowed China’s economy, provincial and municipal governments in many cities across China issued digital coupons to stimulate the economy. The coupons were distributed directly to consumers through pre-existing technology platforms such as Alibaba, Meituan, and JingDong in multiple “coupon waves”. The stated aim of the coupon program was to promote consumption at low fiscal cost. Coupons could only be used in their specific categories to support the recovery of the sectors that local policymakers perceived to have been hit hardest by the pandemic, such as restaurants and tourism.

Most importantly for our analysis, all of the coupons had spending thresholds and discount amounts (“Spend at least ¥ X , get ¥ Y off”), and all of the coupons had a short period in which they needed to be used before they expired (“use it or lose it”). Many municipalities continued to offer coupons throughout the 2021–2023 period partly because of the perceived effectiveness of the initial coupon distributions.

2.2 Data

We use data from one of the large online e-commerce platforms that distributed the coupons. The platform has substantial market share in many different spending categories including restaurants, entertainment, and food delivery.³ In 2018, the platform had more than 600 million registered users and approximately 35 million daily users. We study coupons issued by the platform in three cities.⁴

For each transaction, we observe the spending amount, spending category, and transaction time and date. We merge the transactions data with the platform’s coupon database, which records when the coupon was acquired, the coupon’s threshold and discount, and whether or not the coupon was redeemed. We received data covering all transactions on the platform for three months before and after the coupons were distributed for every consumer who received a coupon during our sample period. To create the data set for analysis, we define the period of each coupon as the number of days each consumer had to use the coupon before it expired. We make sure to include the same days of the week as in 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 as the Tuesday to Saturday of the previous week.

The Appendix gives more details about the data set and the the coupon characteristics. Table OA.1 presents summary statistics for each of the coupons, including the total number of coupons available,

³The data were provided by the platform under a data use agreement that requires us to preserve the anonymity of the platform and the three cities that we focus on. The platform reviewed the study prior to public dissemination for factual inaccuracies, confidential information, and trade secrets.

⁴We report results for all of the coupons in our data except for the coupons distributed in the first wave of coupons in City A. We exclude these coupons because spending changes during that wave are confounded by the 2021 Spring Festival (Lunar New Year); see Appendix A.2 for more details.

the take-up rate, and the redemption rate. Figure OA.3 shows the range of coupon thresholds and discounts in our data. The discounts are always set between 25 and 50 percent of the coupon threshold, implying that when cities chose to offer coupons with higher thresholds, they chose higher discounts, as well.⁵

3 Empirical Approach

3.1 Estimating MPC^{coupon} Using a Bunching Estimator

To estimate the effects of the coupons on spending, we use a bunching estimator that uses the distribution of spending in the period before the coupons were distributed as the counterfactual, following Best et al. (2020), Defusco et al. (2020), and Cengiz et al. (2019). Our bunching estimator takes as an input the distribution of spending in ¥1 bins in the two time periods, the pre-period and coupon-wave period. We estimate the effect of each coupon on spending by calculating the “excess mass” (EM) of transactions above the coupon threshold (τ) and the “missing mass” (MM) of transactions below the coupon threshold using the following bunching estimators:

$$\widehat{EM}_\tau = \sum_{j=\tau}^H (n_j^{WAVE} - n_j^{PRE})j$$

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

where τ denotes the coupon-specific spending threshold, H is a standard tuning parameter that defines the upper bound of the “bunching window”, and n_j^{PRE} and n_j^{WAVE} are the number of transactions with spending amounts between j and $j+1$ yuan in the pre-period and the wave period, respectively.⁶

The sum of the excess mass and missing mass estimates, $\widehat{EM}_\tau + \widehat{MM}_\tau$, is the total effect of the coupons on spending. We define MPC_τ^{coupon} as the increase in spending divided by the total spending by the government:

$$MPC_\tau^{coupon} = \frac{\widehat{EM}_\tau + \widehat{MM}_\tau}{S_\tau} \tag{1}$$

where S_τ is the total government spending on coupons with threshold τ , which equals the per-coupon subsidy τ times the number of coupons redeemed during the coupon wave.

⁵In the Appendix, we describe structured interviews with employees of the platform, who described the municipalities as targeting a “leverage ratio,” which they defined as the ratio of the coupon threshold to the coupon discount amount. Interestingly, this ratio is quite similar to—though not quite the same as—the expression for the MPC^{coupon} that we derive in Section 5 below.

⁶In our main analysis, we set $H = \bar{\tau} + 50$, where $\bar{\tau}$ is the highest coupon threshold across all of the coupons distributed in a given city and spending category.

3.2 Estimating MPC^{coupon} Using Random Assignment

The coupons distributed in one city were randomly assigned within a spending category: conditional on the consumer’s acquisition of a coupon, the threshold and discount were chosen randomly from a set of three options. As a result, we can estimate the causal effect of a consumer’s being assigned the coupon with threshold τ relative to that of being assigned the coupon with threshold τ' by comparing the distribution of spending across the different coupons; there is no need to use pre-period data. We define this causal effect as $MPC_{\tau-\tau'}^{coupon}$ and estimate it as follows:

$$MPC_{\tau-\tau'}^{coupon} = \frac{\sum_{j=1}^H \left[\theta n_{j,\tau}^{WAVE} - (1-\theta)n_{j,\tau'}^{WAVE} \right] j}{\theta S_{\tau} - (1-\theta)S_{\tau'}} \quad (2)$$

where $\theta = Inventory_{\tau'} / (Inventory_{\tau} + Inventory_{\tau'})$ is the share of coupons with threshold τ' . We prove in the Appendix that the coupon-specific bunching estimates from Section 3.1 are related to $MPC_{\tau-\tau'}^{coupon}$ by the following identity:

$$E[MPC_{\tau-\tau'}^{coupon}] = \frac{\theta S_{\tau}}{\theta S_{\tau} - (1-\theta)S_{\tau'}} MPC_{\tau}^{coupon} - \frac{(1-\theta)S_{\tau'}}{\theta S_{\tau} - (1-\theta)S_{\tau'}} MPC_{\tau'}^{coupon} \quad (3)$$

This identity states that the MPC^{coupon} estimated by comparing pairs of randomly assigned coupons is equal to an appropriately-weighted average of the individual MPC^{coupon} estimates recovered from the bunching estimators. A useful implication of this result is that if two coupons have similar MPC^{coupon} estimates, then the government can increase spending by assigning a greater share coupons to the coupon with the higher threshold and discount.

4 Main Results

4.1 Graphical Evidence

We begin by presenting visual evidence of bunching at coupon-specific thresholds. Recall that our data set covers all consumers who acquired coupons, tracking all of their spending on the platform before and after the coupons were distributed.

As a running example, we focus on the 54–18 coupon distributed to City A residents in the second coupon wave. Panel (a) of Figure 1 shows the transaction-level spending distribution in the targeted spending category for recipients of this coupon in the two periods before the coupons were distributed. The similarity between the two pre-period distributions provides evidence against confounding trends in overall spending in the periods before the coupons were distributed.

Next, Panel (b) shows the spending distribution in the coupon-wave period (t) relative to that in the pre-period ($t-1$). This figure shows clear visual evidence of bunching at the coupon-specific threshold. Moreover, to the left of the coupon-specific threshold, there is some visual evidence of “missing mass”, which implies that some consumers spent more than they otherwise would have to be

able to redeem the coupon and earn the discount.⁷ Panel (c) compares the spending distributions for the pre-period ($t - 1$) and the period following the coupon wave ($t + 1$); the distributions are fairly similar, with perhaps some evidence of slightly fewer transactions across the distribution, which would be consistent with a very small amount of intertemporal substitution. Lastly, Panel (d) 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 the choice of pre-period.

The Appendix reports analogous figures for all of the other coupons in our data, and the same patterns consistently emerge: clear visual evidence of bunching at the coupon thresholds, excess mass that is much larger than the missing mass, and no differences in mass in the excluded region in the upper tail (Figures OA.4–OA.17).

4.2 Empirical Estimates of MPC^{coupon}

To quantify the spending effects of the coupons, we estimate equation (1) for each coupon and report bootstrap standard errors for each MPC^{coupon} estimate.⁸ The results are reported in Table 1, which shows that the estimated MPC s range from 1.9 to 4.6, with a weighted average of 3.1–3.3. We immediately evaluate two explanations for these large MPC^{coupon} estimates: substitution between spending categories and intertemporal substitution.

4.2.1 Substitution Between Spending Categories

Since we observe all of the spending on the platform for all of the consumers in our sample, we can estimate the MPC^{coupon} for spending in the non-targeted spending categories. For a supermarket coupon, we can look for evidence of substitution away from spending on other categories such as food delivery, restaurants, and entertainment spending. In Table OA.2, we find no evidence of any statistically or economically significant effects of coupons on the spending in non-targeted spending categories, and in column (6) of Table 1, we find similar MPC^{coupon} estimates when we look at total platform spending. These results suggest that the coupons cause limited substitution between spending categories.

4.2.2 Intertemporal Substitution

To assess the role of short-run intertemporal substitution, we re-estimate equation (1) for multiple additional periods before and after the coupons were distributed, always comparing spending to spending in the $t - 1$ pre-period. Figure OA.18 presents these results, which show no evidence of substantial

⁷Since our analysis uses all transactions made by coupon recipients, the transactions observed immediately to the left of the coupon-specific thresholds do not necessarily indicate that consumers are making dominated choices since they may have used the coupon in a previous transaction during the same period. In Appendix A.3, we investigate this issue in more detail and conclude that dominated choices are infrequent in our setting.

⁸We calculate the bootstrap standard errors based on 1000 replications, using a cluster-based bootstrap procedure that resamples the ¥1 bins of transactions with replacement. In each bootstrap step, we calculate the MPC^{coupon} estimate using equation (1).

intertemporal substitution. As expected from the results in Panel (b) of Figure 1, there is a very small decrease in spending in the $t + 1$ period, but it only offsets the initial increase in spending in period t by only a very small amount.

4.2.3 Robustness and Heterogeneity

We assess the robustness of our main results in two main ways. First, we report similar results when we re-estimate equation (1) using different pre-periods and different values of H , the tuning parameter in the bunching estimation (Table OA.6).

Second, we estimate the MPC^{coupon} using the coupons that were randomly assigned. Panel (a) of Figure 2 shows extremely similar pre-period spending distributions across the consumers assigned different coupons, supporting the validity of the random assignment. Panel (b) of Figure 2 shows that the sharp bunching during the coupon wave lines up exactly with the coupon thresholds assigned to each group of consumers. Using equation (3), we show in Table OA.3 that the estimates based on strict random assignment are always very close to the implied estimates from the bunching estimators.

Lastly, we explore heterogeneity across consumers. We divide consumers into two approximately equal-sized age groups (above and below age 35) and find similar MPC^{coupon} estimates (Table OA.4). We also divide consumers based on how often they used the platform prior to the coupon wave. Somewhat mechanically, the MPC^{coupon} estimates are a bit higher for users not active on the platform, but the results for active users are similar to our baseline estimates (Table OA.5). We also find broadly similar MPC^{coupon} estimates for the most frequent users of the platform, whom we define as consumers who spent regularly across multiple categories. Since we measure only spending on the platform, it is not possible to completely rule out unmeasured “online–offline” substitution, but the similarity in results for the “frequent users” subsample leads us to conclude that online–offline substitution is small.

Overall, our results consistently point toward large MPC^{coupon} estimates that do not come primarily from reduced spending in other categories or from short-run intertemporal substitution. Why then are the MPC^{coupon} estimates so large? The next section develops a simple dynamic model of consumer spending to understand the economics behind the large MPC^{coupon} estimates.

5 Model and Calibration

5.1 Reassessing the Simple Economics of Notches vs. Subsidies

In an early paper on tax notches, [Blinder and Rosen \(1985\)](#) describe a government that tries 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 the simple economics of notches vs. linear subsidies.

Panel (a) of Figure 3 shows a consumer allocating spending between goods A and B and choosing c_A^* and c_B^* . When the government introduces a linear subsidy (τ) on good A , this reduces the price

from p to $p(1 - \tau)$ and rotates out the consumer's budget constraint, leading 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 . [Blinder and Rosen \(1985\)](#) point out that the government could instead design a notch-based incentive where the government transfers an amount ON in cash if the consumer chooses a level of consumption in sector A at or above the notch set at c'_A . The authors then note:

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](#), p737)

We show using the same graphical model that this reasoning is inaccurate. The simple explanation is that, while [Blinder and Rosen's \(1985\)](#) argument that a notch can always be designed to exactly replicate a linear subsidy is correct, the converse does not hold. In particular, the government can design a notch incentive that cannot be exactly replicated by a linear subsidy because the same increase in consumption in sector A would not come at the same revenue cost and would not have the same effect on consumer welfare.

To demonstrate this, Panel (c) shows the government holding constant the cash transfer ON but increasing the notch. The government can continue to increase the notch up to point c''_A , where the consumer is indifferent between increasing consumption up to the notch and receiving cash ON and staying at (c^*_A, c^*_B) .

Finally, Panel (d) shows the linear subsidy that the government would need to choose to achieve the same increase in consumption from c^*_A to c''_A . Not only is this subsidy costlier to the government than the notch incentive, but also the consumer strictly prefers the subsidized outcome to the initial endowment, while the notch policy is designed to increase consumption in sector A with no change to consumer welfare.

These figures illustrate that the government cannot replicate every notch policy with a linear subsidy at the same fiscal cost.⁹ This highlights a 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 a notch to a linear subsidy. In the next section, we build on these graphical results by developing and calibrating a dynamic consumption model to interpret our results.

⁹It is perhaps not surprising that two parameters can be used to replicate any (one-parameter) linear subsidy but there can be two-parameter notches that cannot be exactly replicated by a linear subsidy. In fact, if we combine the linear subsidy with a lump-sum tax, then we can immediately "fix" the claim in [Blinder and Rosen \(1985\)](#) and restore full equivalence. We can illustrate this by vertically shifting down the τ'' line in Panel (d) of [Figure 3](#) so that it intersects with the notch point.

5.2 Consumption Model

5.2.1 Setup

The model is a T -period model with perfect foresight, no uncertainty, and exogenous income.¹⁰ Consumers 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:

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

where $\sigma \equiv 1/(1-\rho)$ 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 determines the share of spending allocated to each sector.

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

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

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

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

where δ is the consumer's subjective discount rate, r is the exogenous interest rate, and y_t is the consumer's exogenous income in each period.

5.2.2 MPC^{coupon} vs. MPC^{cash}

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} as the change in consumption in period 1 relative to the change in income:

$$MPC^{cash} = \frac{\Delta(c_1^A + c_1^B)}{\Delta(y_1)} = \frac{1}{\sum_{t=1}^T \left[(1+r)^{\frac{1-\gamma}{\gamma}} (1+\delta)^{\frac{-1}{\gamma}} \right]^{t-1}}$$

Now consider the government offering a coupon that pays $\text{¥}d$ if the consumer spends more than $\text{¥}D$ in sector A in period 1. We assume that the consumer takes up the coupon if and only if it increases their utility. If the consumer takes up the coupon, then we can define $MPC^{coupon} = \Delta(c_1^A + c_1^B)/d$.

Define c_1^{A*} as the optimal consumption in sector A in period 1 in the absence of a coupon. We cannot solve for MPC^{coupon} analytically, but if $D \leq c_1^{A*}$, then $MPC^{coupon} = MPC^{cash}$ since in this case the coupon is fungible with cash. If $D > c_1^{A*}$, then $MPC^{coupon} > MPC^{cash}$ if the consumer takes

¹⁰For a discussion of recent models that incorporate uncertainty, liquid and illiquid assets, and liquidity constraints, see Kaplan and Violante (2022).

up the coupon, and in this case MPC^{coupon} can be defined as follows:

$$MPC^{coupon} = \frac{D - c_1^{A*}}{d} + \frac{\Delta(c_1^B)}{d} \quad (4)$$

This expression shows that if $\Delta(c_1^B) \approx 0$, then $MPC^{coupon} \approx (D - c_1^{A*})/d$, which is increasing in the coupon threshold and decreasing in the coupon discount. The policymaker can therefore maximize the “bang for the buck” of the coupon by maximizing MPC^{coupon} subject to the constraint that the consumer prefers to take up the coupon.

We can also use the model to calculate the approximate change in utility from receiving a coupon compared to the change in utility from receiving the equivalent amount from the government in cash:

$$\frac{\Delta U^{coupon}}{\Delta U^{cash}} \approx 1 - 0.5 * (1 - \rho) \frac{(\Delta c_1^A)^2}{d * c_1^{A*}} \quad (5)$$

This formula is derived in the Appendix by taking a second-order approximation around the consumer’s utility after receiving d in cash and then “forcing” the consumer to bunch at the coupon threshold. The derivation uses the envelope theorem to ignore all other consumption changes other than Δc_1^A . The quadratic term comes from the second-order approximation and is scaled by $(1 - \rho)$; intuitively, if consumers are very willing to substitute consumption between sectors, then they value the coupon almost as much as cash.

5.2.3 Calibration

We now calibrate our model to illustrate how it can replicate our MPC^{coupon} estimates quantitatively. The calibration parameters are described in the figure notes and are set such that the consumer chooses to spend two percent of their income in the targeted spending category in each period. We also normalize the parameters so that consumption in sector A in period 1 is equal to one in the absence of a coupon. We solve for MPC^{coupon} numerically, and Figure 4 shows how the MPC^{coupon} varies as the coupon threshold rises from $D = 0$ to $D = 3$ (i.e., a threshold equal to three times the spending amount the consumer would have chosen without a coupon), holding the coupon discount constant at $d = 0.3$ throughout.

Figure 4 shows that if the threshold is set below one, then MPC^{coupon} is equal to MPC^{cash} , as expected given fungibility. As the threshold increases from $D = 1$ to $D = 3$, the MPC^{coupon} increases approximately linearly until the coupon’s threshold is high enough that the consumer would experience a decrease in utility from using the coupon. This figure also shows the change in the consumer’s utility from taking up the coupon relative to cash; this change has an inverse-U shape, as expected given the quadratic approximation formula above. For the range of our weighted-average MPC^{coupon} estimates (3.1–3.3), the calibration results indicate that coupons increase consumer utility by approximately 50 percent as much as an equivalent amount of cash. These results can be used to simulated alternative coupon designs. For example, the calibration results show that higher coupon thresholds and discounts

can deliver greater aggregate stimulus, as long as consumers continue to prefer taking up the coupons.

In the Appendix, we present several additional results from the model calibrations, which we briefly summarize here. First, we explore sensitivity to different parameters. We find that the MPC^{coupon} is lower if consumers are more willing to substitute between sectors than over time (Figure OA.20), and that the welfare cost of coupons relative to cash is smaller if consumers are more willing to substitute between sectors. Second, we show that the quadratic approximation formula is very accurate (Figure OA.21), suggesting that the coupon characteristics and ρ are sufficient statistics for analyzing the effects of the coupons on consumer welfare. Lastly, we compare coupons to a temporary tax subsidy that introduces a subsidy τ_A in period 1 but not in any other periods (Figure OA.22). Following Blinder and Rosen (1985), we restrict ourselves to a linear subsidy and compare our coupon to this alternative policy instrument. The calibrations show the potential for notch-based incentives to be strictly preferable to cash transfers and temporary tax subsidies whenever the policymaker puts strong weight on stimulating spending in sector A relative to consumer welfare.

6 Conclusion

This paper studies a novel form of economic stimulus: government-issued digital coupons targeted at specific sectors. Such coupons were distributed across several provinces and municipalities in China in the aftermath of the COVID-19 recession, and the coupons have become popular and continue to be distributed in many cities in China. Using data from a large e-commerce platform, we estimate large effects of the coupons on spending. We rule out cross-category and intertemporal substitution as the primary explanations for our large spending estimates, and we develop a dynamic model that rationalizes the large MPC^{coupon} estimates as arising from the temporary notches created by the coupons.

If policymakers are primarily interested in supporting targeted sectors, then our model makes clear why coupons can have attractive targeting properties. In the model, 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 immediately increase spending in the *targeted* sectors to receive the coupon discount. Tax notches are often seen as a “design flaw” in public finance since it is difficult to imagine an optimal tax policy featuring a tax notch. When it comes to fiscal stimulus, however, the incentives created by the digital coupons may be a feature rather than a bug.

Given the novelty of this type of stimulus, we see several areas for future work. First, our analysis abstracted from many types of consumer heterogeneity. While our heterogeneity analysis found broadly similar MPC^{coupon} estimates by age and prior activity on the platform, we know from Blinder and Rosen (1985) that heterogeneity in behavioral responses to notches is a key factor in determining the attractiveness of notches compared to linear subsidies.

Second, we discussed differences between the effects of coupons and the effects of cash transfers, but we did not find existing MPC^{cash} estimates for Chinese consumers to benchmark against our

MPC^{coupon} estimates. Future work should produce MPC^{cash} estimates specific to China, perhaps by using the kind of natural experiments surveyed by [Kaplan and Violante \(2022\)](#) or by carrying out a randomized cash transfer experiment as in [Boehm et al. \(2023\)](#).

Finally, our model-based analysis focused primarily on understanding the MPC^{coupon} estimates, but consumers also decide whether to take up and use the coupon. Our model shows that if the coupon threshold is set “too high,” many consumers will not use the coupon. Additionally, we observe in the data that many of the coupons that were taken up were not used. Incomplete take-up and incomplete redemption reduce the aggregate impact of coupons, and future work should model the additional trade-offs that come from consumers’ take-up and redemption decisions. These theoretical and empirical extensions should help provide policymakers with additional information to guide the optimal design of targeted fiscal stimulus using digital coupons.

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 Ilzetzi, 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.
- Boehm, Johannes, Etienne Fize, and Xavier Jaravel**, “Five Facts about MPCs: Evidence from a Randomized Experiment,” *Sciences Po Economics Discussion Paper No. 2023-02*, 2023.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer**, “The Effect of Minimum Wages on Low-wage Jobs,” *Quarterly Journal of Economics*, 2019.
- Correia, Isabel, Emmanuel Farhi, Juan Pablo Nicolini, and Pedro Teles**, “Unconventional Fiscal Policy at the Zero Bound,” *American Economic Review*, June 2013, *103* (4), 1172–1211.
- , **Juan Pablo Nicolini, and Pedro Teles**, “Optimal Fiscal and Monetary Policy: Equivalence Results,” *Journal of Political Economy*, 2008, *116* (1), 141–170.
- Defusco, Anthony, Stephanie Johnson, and John Mondragon**, “Regulating Household Leverage,” *Review of Economic Studies*, 2020, *87*.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and IvÅ;n Werning**, “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?,” *American Economic Review*, May 2022, *112* (5), 1437–74.
- Hsieh, Chang-Tai, Satoshi Shimizutani, and Masahiro Hori**, “Did Japan’s shopping coupon program increase spending?,” *Journal of Public Economics*, 2010, *94*, 523–529.
- 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.
- Kaplan, Greg and Giovanni L Violante**, “The marginal propensity to consume in heterogeneous agent models,” *Annual Review of Economics*, 2022, *14*, 747–775.
- Kleven, Henrik**, “Bunching,” *Annual Review of Economics*, 2016, *8*, 435–464.
- and **Mazhar Waseem**, “Using Notches to Uncover Optimization Frictions and Structural Elasticities: 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.

Shapiro, Matthew D and Joel Slemrod, “Did the 2008 tax rebates stimulate spending?,” *American Economic Review*, 2009, 99 (2), 374–79.

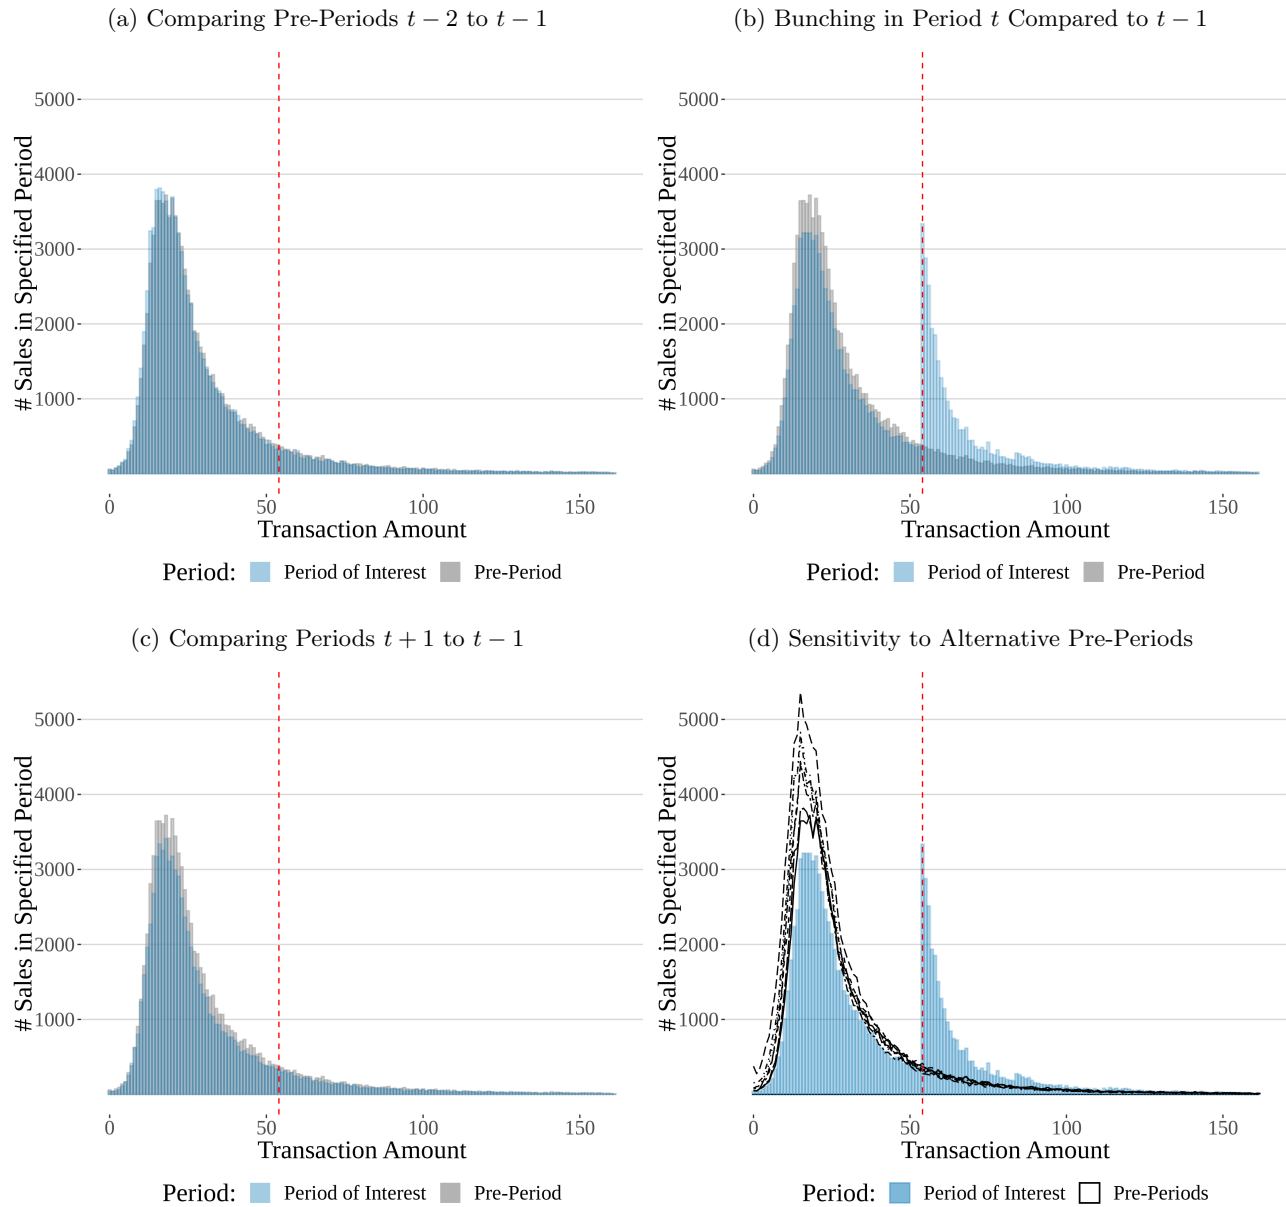
Xing, Jianwei, Eric Zou, Zhentao Yin, Yong Wang, and Zhenhua Li, “‘Quick Response’ Economic Stimulus: The Effect of Small-Value Digital Coupons on Spending,” *NBER Working Paper No. 27596*, 2021.

Table 1
Bunching Estimates of Effects of Coupons on Spending

City	Spending Category	Coupon Wave	Coupon [Threshold-Discount]	MPC^{coupon}	
				Spending in Targeted Category	Total Spending on Platform
(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Coupon-Specific MPC^{coupon} Estimates					
City A	Supermarket	2	24-8	3.94 (0.16)	4.59 (0.39)
City A	Supermarket	2	54-18	3.82 (0.07)	4.10 (0.24)
City A	Supermarket	2	84-28	3.50 (0.04)	3.62 (0.28)
City A	Multi-Category	2	54-18	3.05 (0.14)	3.10 (0.14)
City A	Multi-Category	2	84-28	2.82 (0.15)	2.89 (0.15)
City A	Multi-Category	2	114-38	2.37 (0.18)	2.42 (0.19)
City B	Food Delivery	1	30-15	2.56 (0.16)	2.65 (0.35)
City B	Food Delivery	2	30-15	1.96 (0.25)	2.13 (0.29)
City C	Multi-Category	1	100-40	3.33 (0.07)	3.31 (0.07)
City C	Multi-Category	1	200-100	1.91 (0.14)	1.90 (0.15)
City C	Multi-Category	2	100-40	3.26 (0.09)	3.29 (0.09)
City C	Multi-Category	2	200-100	1.93 (0.15)	1.94 (0.16)
Panel B: Weighted-Average MPC^{coupon} Estimates					
Weight by Number of Coupons Distributed				3.13	3.28
Weight by Number of Coupons Taken Up				3.15	3.31
Weight by Number of Coupons Redeemed				3.11	3.20

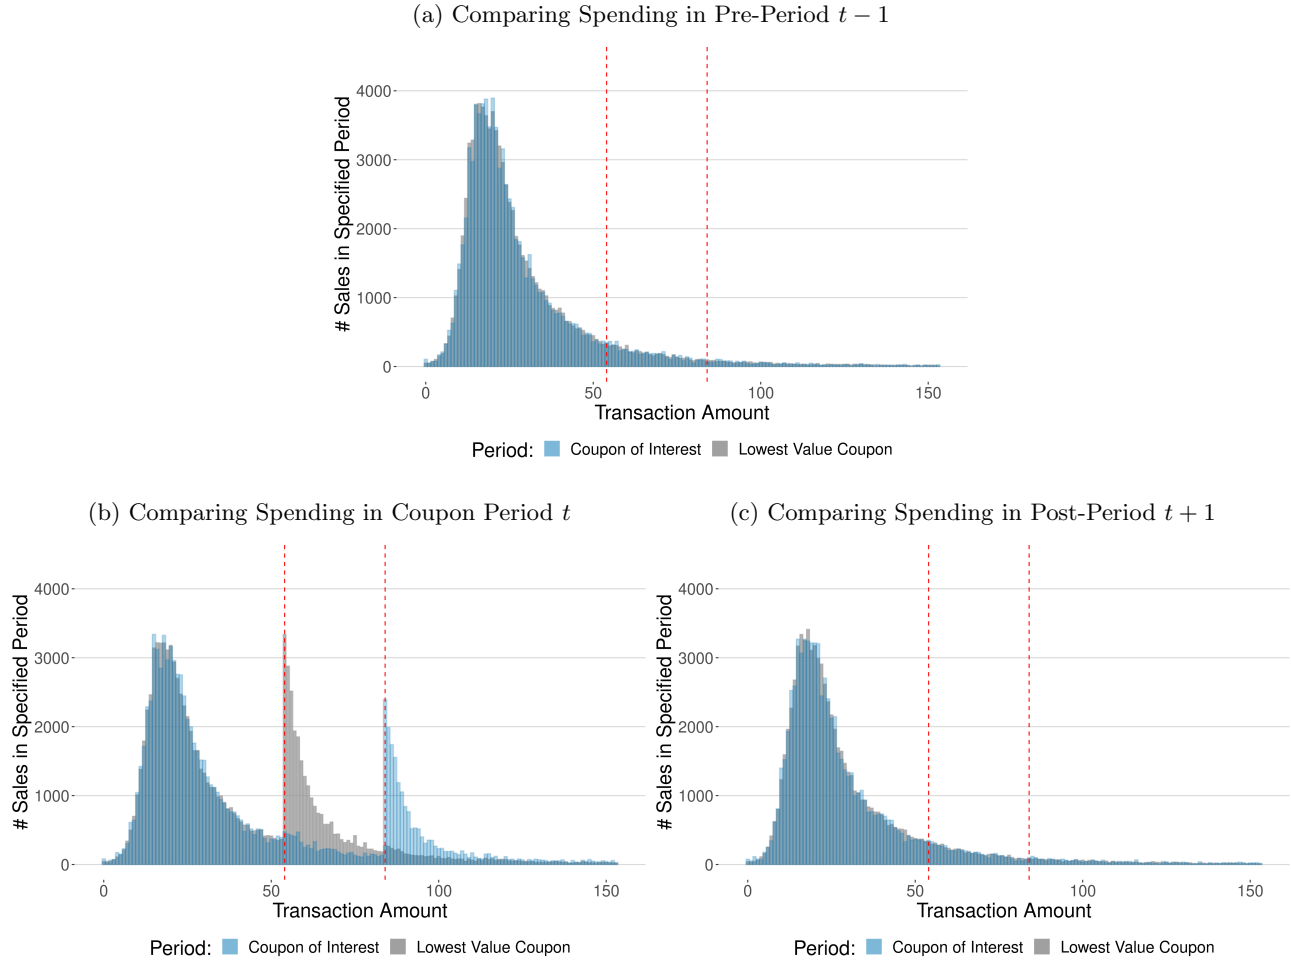
Notes: This table presents coupon MPC estimates using the bunching estimator described in equation (1). Column (1) reports the anonymized city the coupon was distributed in, and columns (2) through (4) describe additional details of the coupon. Column (5) reports the coupon MPC estimate within the targeted spending category. Column (6) reports the coupon MPC estimate for total spending. Bootstrap standard errors are presented in parentheses, based on 1000 replications of a cluster-based bootstrap procedure that resamples the ¥1 bins of transactions with replacement.

Figure 1
Illustration of Bunching Estimator for 54-18 Coupon in City A



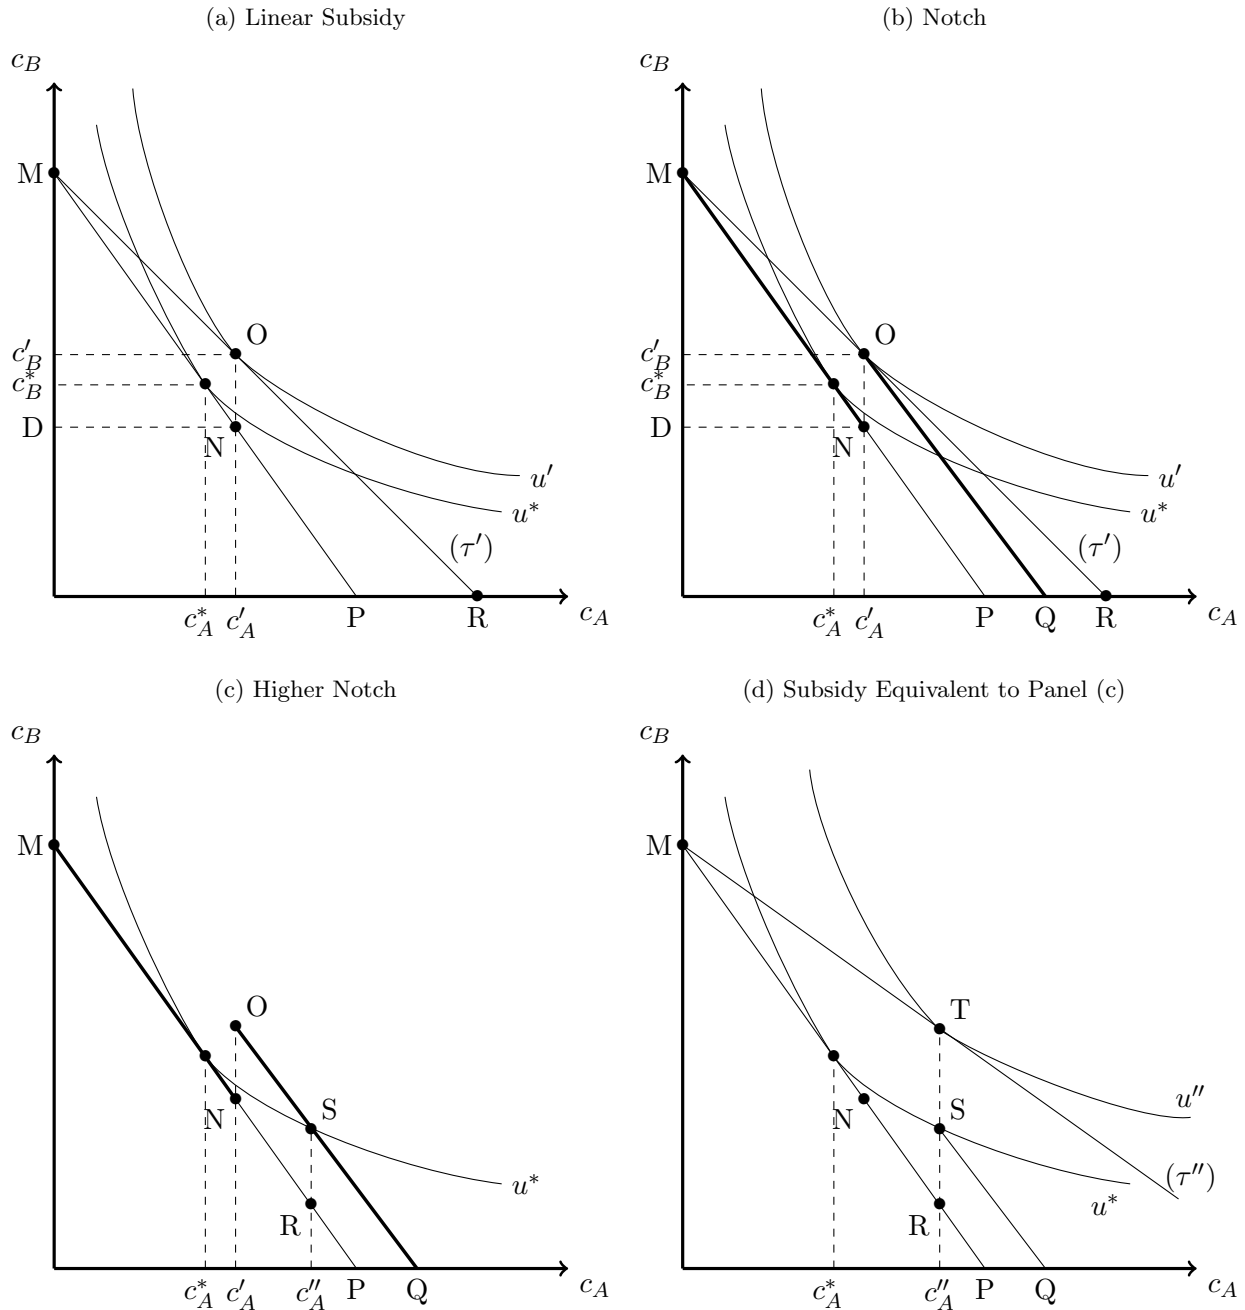
Notes: This figure illustrates the bunching estimator by comparing the distribution of food delivery spending between periods around the time the coupons were distributed. Panel (a) compares the distribution of spending in the two pre-periods immediately before the coupons were distributed. Panel (b) 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 panels (a) to (c) the pre-period $t - 1$ distribution is shown for reference. Panel (d) illustrates the sensitivity to different pre-periods by comparing the distribution in the coupon period to seven pre-periods ($t - 1$ through $t - 7$). The analogous figure covering all of the spending categories covered by the coupon is available in the Appendix (see Figure OA.15).

Figure 2
 Estimating MPC^{coupon} Using the Random Assignment of Coupons



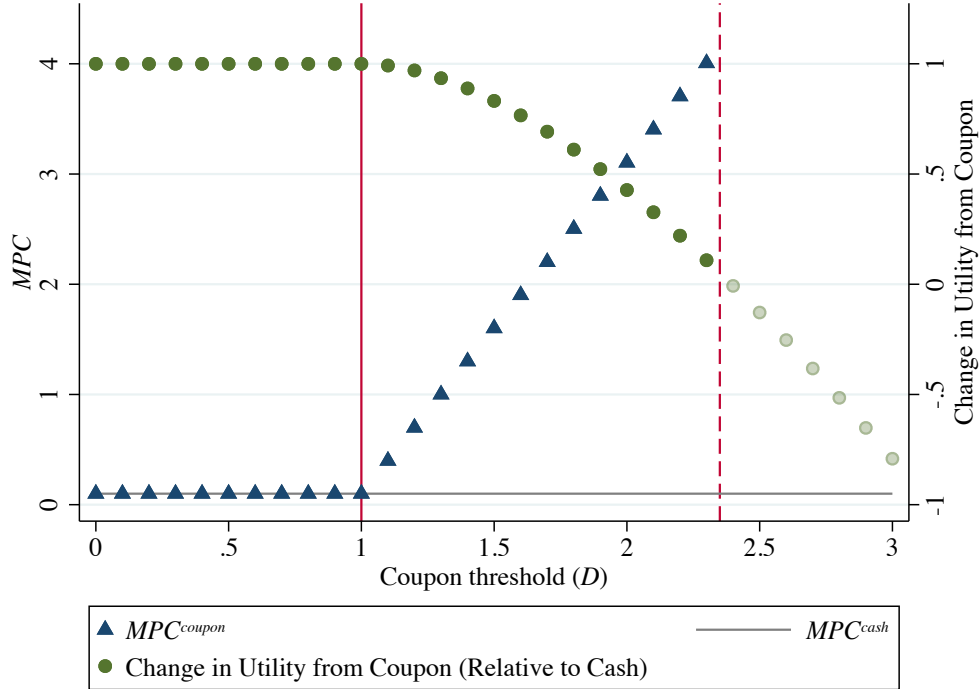
Notes: This figure reports panels analogous to Figure 1 except that the identification is based on comparing the consumers who were randomly assigned different coupons in Wave 2 in City A. Panel (a) compares the distribution of spending between the two groups of consumers assigned either the 54-18 or 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 period does not show up as lower spending in the following period.

Figure 3
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 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 4
Model Calibration



Notes: This figure shows how MPC^{coupon} varies with the coupon threshold (D) and compares MPC^{coupon} to MPC^{cash} . The model calibration uses the following parameters: $T = 10$, $y_t = 50$ in each period, $r = \delta = 0$ (i.e., no discounting and no borrowing costs), $\gamma = 0.5$, $\rho = 0.5$, and $\alpha = 0.125$ (so that the consumer chooses to spend 2 percent of their income in sector A each period in the absence of a coupon; i.e., $c_t^A = 1$). The coupon discount is held fixed at $d = 0.3$ as D varies. When the coupon threshold is less than 1, the coupon is fungible with cash, which implies $MPC^{coupon} = MPC^{cash}$. As the coupon threshold continues to increase above 1, MPC^{coupon} increases, and it increases approximately linearly as would be expected based on equation (4) when $\Delta(c_1^B) \approx 0$ (which happens to be the case at these parameter values). The solid circles show that as the coupon threshold increases above 1, the increase in consumer utility from using the coupon decreases relative to the increase in consumer utility from an equivalent amount of government spending distributed as cash. Once the coupon threshold crosses the dashed vertical line, the consumer is worse off using the coupon, and so MPC^{coupon} is no longer defined because the consumer would not choose to use the coupon.