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

Jing Ding — Tongji University

Lei Jiang — Kent State University

Lucy Msall — University of Chicago

Matthew J. Notowidigdo — University of Chicago and NBER

First Version: February 25, 2022

This Version: August 26, 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 Connor Fogal, I-Han Lai, Angelo Marino, Jeremy Mopsick, and Ningxin Zhang for excellent research assistance and Milena Almagro, Peter Ganong, Zhiguo He, Anders Humlum, Koichiro Ito, Emir Kamenica, Kory Kroft, Maria Polyakova, Nancy Qian, Eric Zou, Eric Zwick, and numerous seminar participants for helpful comments. We also thank Sam Eyler-Driscoll for helpful editing. 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 governments 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 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 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, the Chinese government carried out a novel form of stimulus using publicly-financed digital coupons. The coupons were delivered through smartphone apps and designed to stimulate spending in targeted sectors such as restaurants, grocery stores, and shopping malls, since these sectors were hit particularly hard during the early months of the COVID-19 pandemic. The coupons had fixed spending thresholds that needed to be reached before consumers received money from the government—for example, a coupon would give ¥18 off a food delivery order if the total transaction amount was at least ¥54 ("Spend at least ¥54, get ¥18 off").

In this paper, we estimate the effects of the 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 each transaction on the platform for everyone who received coupons. The coupons had a range of different spending thresholds and applied to different spending categories. Throughout the paper, we define the 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 are 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 below, the reason 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 coupons 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 spending thresholds. We 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 identify and estimate the MPC^{coupon} coupon by coupon by integrating over the difference in spending distributions between periods.

We find a range of MPC^{coupon} estimates across the coupons (1.9 to 4.6), with a weighted average of 3.1–3.3. We find no evidence of meaningful substitution between "targeted" and "non-targeted" spending categories using data on all of the consumers' spending on the platform. 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. Lastly, we find very similar results from an alternative empirical approach that exploits the random assignment of coupons for a subset of the coupons in our data. As far as we know, this is the first time bunching estimates are validated using explicit random assignment.

In the final part of the paper, we develop a simple dynamic model of consumer spending to interpret our reduced-form results. We show that the model can match our MPC^{coupon} estimates if the coupon threshold is set higher than the spending in the targeted sector that consumers would have chosen in the absence of the coupon. We also use the 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.

Taken together, our empirical and theoretical results suggest that digital coupons are a cost-effective way to provide targeted stimulus to specific sectors.[1] Our MPC^{coupon} estimates are large and persist for several months, implying that the increased spending from the coupons is achieved at a very low cost relative to other forms of fiscal 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

[1] Throughout 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 these studies focus on state-specific rather than sector-specific tax instruments. We conjecture that the recent analysis of "Keynesian supply shocks" during a pandemic (Guerrieri et al. 2022) can be extended to provide a more rigorous justification for when the government would want to provide a targeted temporary tax cut to a specific sector. If so, then our analysis suggests that it may be preferable for the government to use temporary notches rather than temporary tax subsidies to cost-effectively increase spending in particular sectors.

(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).

Lastly, our paper is most closely related to two other recent studies of digital coupons in China. 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 consumers who just missed out on receiving a coupon to consumers who 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 these studies, our paper has three key advantages. First, our data is broader in scope. Our data covers a larger number of cities and coupons and a wider range of thresholds and discounts, which makes our estimates more readily generalizable to the coupons distributed around the country.[3] Second, we are the first to use a bunching estimation approach and exploit the explicit random assignment of coupon thresholds and discounts. The random assignment allows us to draw stronger inferences about the causal effect of coupons on spending relative to the other studies.[4] Lastly, we develop a dynamic model to rationalize the large MPC^{coupon} estimates, calculate the consumer welfare effects of coupons relative to cash, and compare coupons to temporary tax subsidies.

2 Background and Data

2.1 Background on the Chinese Coupon Programs

In response to the COVID-19 pandemic that slowed China's economy, local governments across China financed digital coupons to stimulate the economy. The coupons were distributed directly to consumers through pre-existing technology platforms such as Alibaba, JingDong, and Meituan. The stated aim of the coupon program was to stimulate consumption at low fiscal cost. Coupons could only be used in specific categories to support the recovery of the sectors that local policymakers considered to be hit hardest by the pandemic (such as restaurants and tourism). All of the coupons had spending

[2] Xing 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.

[3] Our study also uses data from a different online platform and covers a different set of cities and coupon waves.

[4] An additional methodological strength of our bunching estimator approach is that it can 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 with incomplete take-up, which is the case for 7 of the 12 coupons in our data.

thresholds and discount amounts. Consumers acquired coupons within the platforms' mobile apps by tapping on an icon, and the coupons were typically distributed "first come, first served." Once consumers acquired coupons, they could redeem them on the platform within a short time window before they expired.

2.2 Data

We use data from one of the large online e-commerce platforms that distributed the coupons in 2021. The platform has substantial market share in many different spending categories including restaurants, entertainment, and food delivery.[5] In 2018, the platform had more than 600 million registered users and approximately 35 million daily users, and we study coupons issued by the platform in three cities.[6]

For each transaction, we observe the time, date, spending amount, and spending category. We merge the transactions data with the platform's coupon database, which records when each 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.[7]

The Appendix provides additional 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, the take-up rate, and the redemption rate. Table OA.2 gives the issue date and duration of each coupon. Lastly, Figure OA.3 plots the coupon thresholds and discounts in our data. The discount is always set between 25 and 50 percent of the coupon threshold, implying that when policymakers chose to offer coupons with higher thresholds, they chose higher discounts, as well.[8]

[5] The 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 did not review the study prior to public dissemination.

[6] We 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. Additionally, in 2020 many local governments implemented strict "lockdown" measures throughout the country, but these were relaxed by 2021. There were no city-wide "lockdowns" during our sample period in any of the three cities, although there were occasional lockdowns of buildings within the cities that typically lasted 1-2 weeks. The three cities also continued various social distancing restrictions throughout the year at public places such as shopping malls and restaurants.

[7] 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; e.g., if a coupon was available to use for 5 days from Tuesday to Saturday, then we define the pre-period as the Tuesday to Saturday of the previous week.

[8] In 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's threshold to the discount amount. Interestingly, this ratio is quite similar to—though not quite the same as—the analytical expression for the MPC^{coupon} that we derive in Section 5.

3 Empirical Approach

3.1 Estimating MPC^{coupon} Using a Bunching Estimator

We estimate the effects of the coupons on spending using a bunching estimator that takes 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). The bunching estimator uses 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 as follows:

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

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

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 wave period, respectively.[9]

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 discount amount τ times the number of coupons redeemed during the coupon wave. We use the same approach to estimate the MPC^{coupon} for spending in the targeted category as well as total spending, depending on the set of transactions used in the estimation.

3.2 Estimating MPC^{coupon} Using Random Assignment

The coupons distributed in city A were randomly assigned within a spending category: conditional on a consumer acquiring 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 being assigned a coupon with threshold τ relative to being assigned a coupon with threshold τ' simply by comparing the distribution of spending across the consumers assigned the different coupons; there is no need to use pre-period

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

data. We define this causal effect as $MPC^{coupon}_{\tau-\tau'}$ and estimate it as follows:

$$MPC^{coupon}_{\tau-\tau'} = \frac{\sum_{j=1}^{H}\left[\theta n^{WAVE}_{j,\tau} - (1-\theta)n^{WAVE}_{j,\tau'}\right]j}{\theta S_{\tau} - (1-\theta)S_{\tau'}} \tag{2}$$

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 the $MPC^{coupon}_{\tau-\tau'}$ estimate by the following identity:

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

This identity states that the $MPC^{coupon}_{\tau-\tau'}$ estimate formed by comparing two randomly assigned coupons is equal to an appropriately-weighted average of the individual MPC^{coupon} estimates recovered from bunching estimators. A useful economic implication of this result is that if two coupons have similar MPC^{coupon} estimates, then the government can increase aggregate spending by assigning a greater share of the 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 spending thresholds. Recall that our data covers all consumers who acquired coupons and includes all of their transactions on the platform before and after the coupons were distributed.

As a running example, we focus on the 54–18 coupon distributed in City A in the second coupon wave. Panel (a) of Figure 1 shows the spending distribution in the coupon-wave period (t) relative to the pre-period ($t-1$). The figure shows clear visual evidence of bunching at the coupon-specific spending threshold. Moreover, to the left of the coupon-specific threshold, there is clear evidence of "missing mass", which implies that the coupons caused some consumers to spend more than they otherwise would have in order to redeem the coupon and earn the discount.[10]

Panel (b) compares the spending distributions in the pre-period ($t-1$) to the period following the coupon wave ($t+1$); the distributions are fairly similar, with slightly fewer transactions in the post-period, which would be consistent with a very small amount of intertemporal substitution. Panel (c) shows that the pre-period spending distributions are fairly stable for several periods leading up to the coupon wave, which provides evidence against confounding trends in spending and indicates that our results are not sensitive to the choice of pre-period. Lastly, Panel (d) compares the consumers

[10] Since 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 in more detail and conclude that dominated choices are infrequent in our setting.

who were randomly assigned different coupons; the figure shows that the sharp bunching during the coupon wave lines up exactly with the coupon-specific thresholds assigned to each group of consumers.

The Appendix reports analogous figures for the other coupons in our data, and the same patterns consistently emerge: clear visual evidence of bunching at 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.6–OA.19).

4.2 Empirical Estimates of MPC^{coupon}

To quantify the spending effects of coupons, we estimate equation (1) for each coupon and report standard errors for each MPC^{coupon} estimate.[11] The results are reported in column (1) in Table 1 and show that the estimated MPC^{coupon} estimates range from 1.9 to 4.6, with a weighted average of 3.1–3.3. Using equation (3), Table 2 shows that the MPC^{coupon} estimates based on random assignment are always very close to the corresponding estimates from the bunching estimators, validating the bunching estimates. We next explore two explanations for the large MPC^{coupon} estimates: substitution between spending categories and intertemporal substitution.

4.2.1 Substitution Between Spending Categories.

Since we observe all spending on the platform for the consumers in our sample, we estimate the MPC^{coupon} for total spending on the platform as well as total spending in the non-targeted spending categories. The results in columns (2) and (3) in Table 1 show that total spending on the platform increases by about the same amount as the increase in the targeted spending categories, and we find no evidence of any statistically or economically significant effects of coupons on total spending in non-targeted spending categories. These results suggest that the coupons cause limited substitution between spending categories.[12]

As described above, we only observe consumer spending on the platform. Using data from the platform's annual reports, we estimate that the platform's users spend ¥1,890 per year on the platform on average. Using survey data from the China Family Panel Studies (CFPS) and the China Household Finance Survey (CHFS), we estimate average annual expenditure between ¥33,092 and ¥34,390, which implies that we capture roughly 5.5–5.7 percent of total consumption.[13] To investigate potential bias from "online-offline" substitution, we divide consumers based on how often they used the platform prior to the coupon distributions. Somewhat mechanically, the MPC^{coupon} estimates are a bit higher

[11] We calculate bootstrap standard errors based on 1000 replications using a cluster-based bootstrap procedure that resamples the ¥1 bins of transactions with replacement.

[12] In Table OA.3, we disaggregate the non-targeted spending into several different sub-categories, and we find suggestive evidence that supermarket coupons lead to a decrease in food delivery spending and an increase in spending on movies, but we interpret these results cautiously since the event study estimates for the subcategories are fairly volatile.

[13] We use 2020 CFPS and 2021 CHFS data restricted to respondents living in urban areas. We re-weight the survey data to match the age distribution of users on the platform and estimate average wage income. We then assume that 67.4 percent of income is spent on consumption based on government reports of per capita income and expenditure (National Bureau Of Statistics Of China 2023).

for users who were not active on the platform before the coupons were distributed, 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, defined as the consumers who spend regularly across multiple categories. Thus, while we cannot completely rule out "online–offline" substitution with our data, the results for the most active users of the platform lead us to conclude that this potential bias is limited in our setting.

4.2.2 Intertemporal Substitution

To assess the role of short-run intertemporal substitution, we re-estimate equation (1) for multiple periods before and after the coupons were distributed, comparing spending in each period to the $t-1$ pre-period. Panel (a) of Figure 3 shows spending in the targeted spending category, and the results show a small decrease in the $t+1$ period, which only offsets the initial increase in spending by a very small amount. Panel (b) shows results for all other spending on the platform, which indicate very little cross-category substitution over time, and Panel (c) shows results for total spending on the platform, which largely matches the results in Panel (a). Lastly, Panel (d) shows a similar pattern of results comparing the total spending on the platform for consumers randomly assigned different coupons. Taken together, the results in Figure 3 indicate very limited intertemporal substitution.

The limited intertemporal substitution in our setting contrasts somewhat with the results in Xing et al. (2021), which show a more important role for intertemporal substitution. A likely explanation for the difference in results is that in our setting the coupons targeted non-durable spending categories such food delivery and supermarkets, while the strongest intertemporal substitution effects in Xing et al. (2021) come from the coupons for books and cellphones, with much more limited intertemporal substitution for the dining coupons. Additionally, our coupons are less generous on average than the coupons in Xing et al. (2021), partly because the cellphone coupons in their setting had very high discounts. As a result, we find an average increase in spending of ¥77.55 compared to ¥299.34 in Xing et al. (2021), which would also explain some of the difference in results if larger coupons tend to lead to more intertemporal substitution.

4.2.3 Robustness and Heterogeneity

We assess the robustness of our main results by re-estimating equation (1) using different values of H (the tuning parameter in the bunching estimation), and we find similar results (Table OA.6). We also explore heterogeneity across consumers by dividing consumers into two roughly equal-sized age groups (above and below age 35), and we find similar MPC^{coupon} estimates across the two groups (Table OA.4).

Overall, we conclude that our large MPC^{coupon} estimates 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 Consumption Model

5.1 Reassessing the Simple Economics of Notches Versus Subsidies

Since the coupon thresholds create temporary notches in the consumer budget constraint, we begin by presenting a simple graphical model to reassess the simple economics of notches versus subsidies. In an early paper on tax notches, Blinder and Rosen (1985) describe a government that stimulates consumption of a given commodity through a linear subsidy, and we adopt their single representative agent framework in this subsection.

Panel (a) of Figure 2 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 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.[14] 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. We now build on these graphical results by developing a dynamic consumption model.

5.2 Model Setup

The model is a T-period model with perfect foresight, no uncertainty, and exogenous income.[15] Consumers borrow, save, and allocate consumption across time periods and sectors.

The consumer's per-period utility function is given by:

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

where $\sigma \equiv 1/(1-\rho)$ is the consumer's elasticity of substitution between consumption in sectors A and B (c_t^A and c_t^B), $1/\gamma$ is the intertemporal elasticity of substitution, and α is a share parameter.

The consumer's lifetime utility function is given by:

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

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_t is the consumer's exogenous income in each period.

5.3 MPC^{coupon} versus 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:

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

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

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

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

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

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

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.

5.4 Consumer Welfare

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:

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

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.

5.5 Model Simulations

In the Appendix, we simulate the model and show that we can choose plausible parameter values to replicate our MPC^{coupon} estimates quantitatively. The simulations also show how the MPC^{coupon} increases as the coupon threshold increases up to the point where the consumer no longer chooses to take up the coupon (Figure OA.21). Up to this point, the change in consumer's utility follows an inverse-U shape as expected given the quadratic approximation formula above (Figure OA.23).[16] We

[16] With the caveat that our model simulations are highly stylized, when we choose parameters to match the MPC^{coupon} estimates, we find that coupons increase consumer welfare by approximately 50 percent as much as an equivalent amount of cash.

also use the simulations to compare coupons to a temporary subsidy by introducing a linear subsidy τ_A only in period 1 (Figure OA.24). Consistent with our graphical model, we find that coupons (with their notch-based incentives) outperform cash transfers and temporary subsidies whenever the policymaker puts substantial weight on stimulating spending in sector A relative to increasing consumer welfare.

6 Conclusion

We study a novel form of fiscal stimulus: publicly-financed digital coupons targeted at specific sectors. Such coupons were distributed across many provinces and municipalities in China in the aftermath of the COVID-19 recession. The coupons quickly became popular and continued to be distributed as the Chinese economy recovered from the pandemic.

Using data from a large e-commerce platform, we estimate large effects of the coupons on spending, and we develop a dynamic model that rationalizes the large MPC^{coupon} estimates as arising from the temporary notches created by the coupons. Our model makes clear why coupons have attractive targeting properties. In the model, cash distributed by the government is mostly spent on *non-targeted* sectors and saved for the future. The time-limited coupons, however, cause consumers to immediately increase spending in the *targeted* sectors. 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 notch incentives created by 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 sources of consumer heterogeneity. While our heterogeneity analysis found fairly similar MPC^{coupon} estimates by age and prior activity on the platform, we know from Blinder and Rosen (1985) that heterogeneity in behavioral responses is a key factor in determining the attractiveness of notch-based incentives compared to linear subsidies.

Second, we compared the effects of coupons to the effects of cash transfers analytically, 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). Additionally, there may be important distributional consequences of coupons as compared to cash. For example, low-income consumers may not have the cash on hand to be able to reach the spending thresholds and take advantage of coupons, which would make coupons a regressive stimulus policy.

Third, we did not have data to measure the effects of the coupons on prices. Since the coupon transactions represent a small share of total transactions on the platform, we do not think that firms adjusted their prices in response to the coupons, but this may become more important if the coupon distributions are scaled up substantially in the future.

Finally, our model-based analysis focused on understanding the large MPC^{coupon} estimates, but consumers also decide whether to take up and use the coupon, and many of the coupons that were taken up were not used. Both incomplete take-up and incomplete redemption reduce the aggregate impact of coupons. Future work should model the additional trade-offs that come from consumers' take-up and redemption decisions. These extensions would help provide policymakers with additional guidance for 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 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.

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 IvÃ¡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 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.

National Bureau Of Statistics Of China, "Households' Income and Consumption Expenditure in 2022," https://www.stats.gov.cn/english/PressRelease/202301/t20230118_1892303.html Jan 2023. Accessed: 2024-08-19.

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	Non-Targeted Spending on Platform
(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Coupon-Specific MPC^{coupon} Estimates						
City A	Supermarket	2	24-8	3.94	4.59	0.65
				(0.16)	(0.39)	(0.36)
City A	Supermarket	2	54-18	3.82	4.10	0.28
				(0.07)	(0.24)	(0.23)
City A	Supermarket	2	84-28	3.50	3.62	0.12
				(0.04)	(0.28)	(0.28)
City A	Multi-Category	2	54-18	3.05	3.10	0.05
				(0.14)	(0.14)	(0.04)
City A	Multi-Category	2	84-28	2.82	2.89	0.07
				(0.15)	(0.15)	(0.05)
City A	Multi-Category	2	114-38	2.37	2.42	0.05
				(0.18)	(0.19)	(0.06)
City B	Food Delivery	1	30-15	2.56	2.65	0.09
				(0.16)	(0.35)	(0.31)
City B	Food Delivery	2	30-15	1.96	2.13	0.17
				(0.25)	(0.29)	(0.15)
City C	Multi-Category	1	100-40	3.33	3.31	-0.02
				(0.07)	(0.07)	(0.02)
City C	Multi-Category	1	200-100	1.91	1.90	-0.01
				(0.14)	(0.15)	(0.07)
City C	Multi-Category	2	100-40	3.26	3.29	0.03
				(0.09)	(0.09)	(0.03)
City C	Multi-Category	2	200-100	1.93	1.94	0.01
				(0.15)	(0.16)	(0.06)
Panel B: Weighted Average MPC^{coupon} Estimates						
Weight by Number of Coupons Distributed				3.13	3.28	0.16
Weight by Number of Coupons Taken Up				3.15	3.31	0.16
Weight by Number of Coupons Redeemed				3.11	3.20	0.09

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.

Table 2
Validating Bunching Estimates Using Random Assignment of Coupons

City	Spending Category	Coupon Wave	Coupons, τ | τ' [Threshold-Discount]	Effect of Coupon τ' on Spending Relative to Coupon τ: Effect Based on Bunching Estimates	Effect of Coupon τ' on Spending Relative to Coupon τ: Effect Based on Random Assignment	Difference Between Bunching and Randomized Estimates	Percent Difference
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A Coupon-Specific MPC^{coupon} Estimates							
City	Supermarket	2	24-8 | 54-18	3.77	3.69	-0.08	-2.3%
City	Supermarket	2	24-8 | 84-28	3.41	3.41	0.00	-0.1%
City	Multi-Category	2	54-18 | 84-28	2.34	2.34	0.00	-0.2%
City	Multi-Category	2	54-18 | 114-38	1.58	1.70	0.12	6.9%

Notes: This table presents coupon *MPC* estimates using random assignment. Column (1) reports the anonymized city the coupon was distributed in, and columns (2) through (4) describe additional details of the coupon. Column (5) is calculated following equation (3) using *MPC* estimates from Table 1 and the number of coupons redeemed from Table OA.1

Figure 1
Illustration of Bunching Estimators

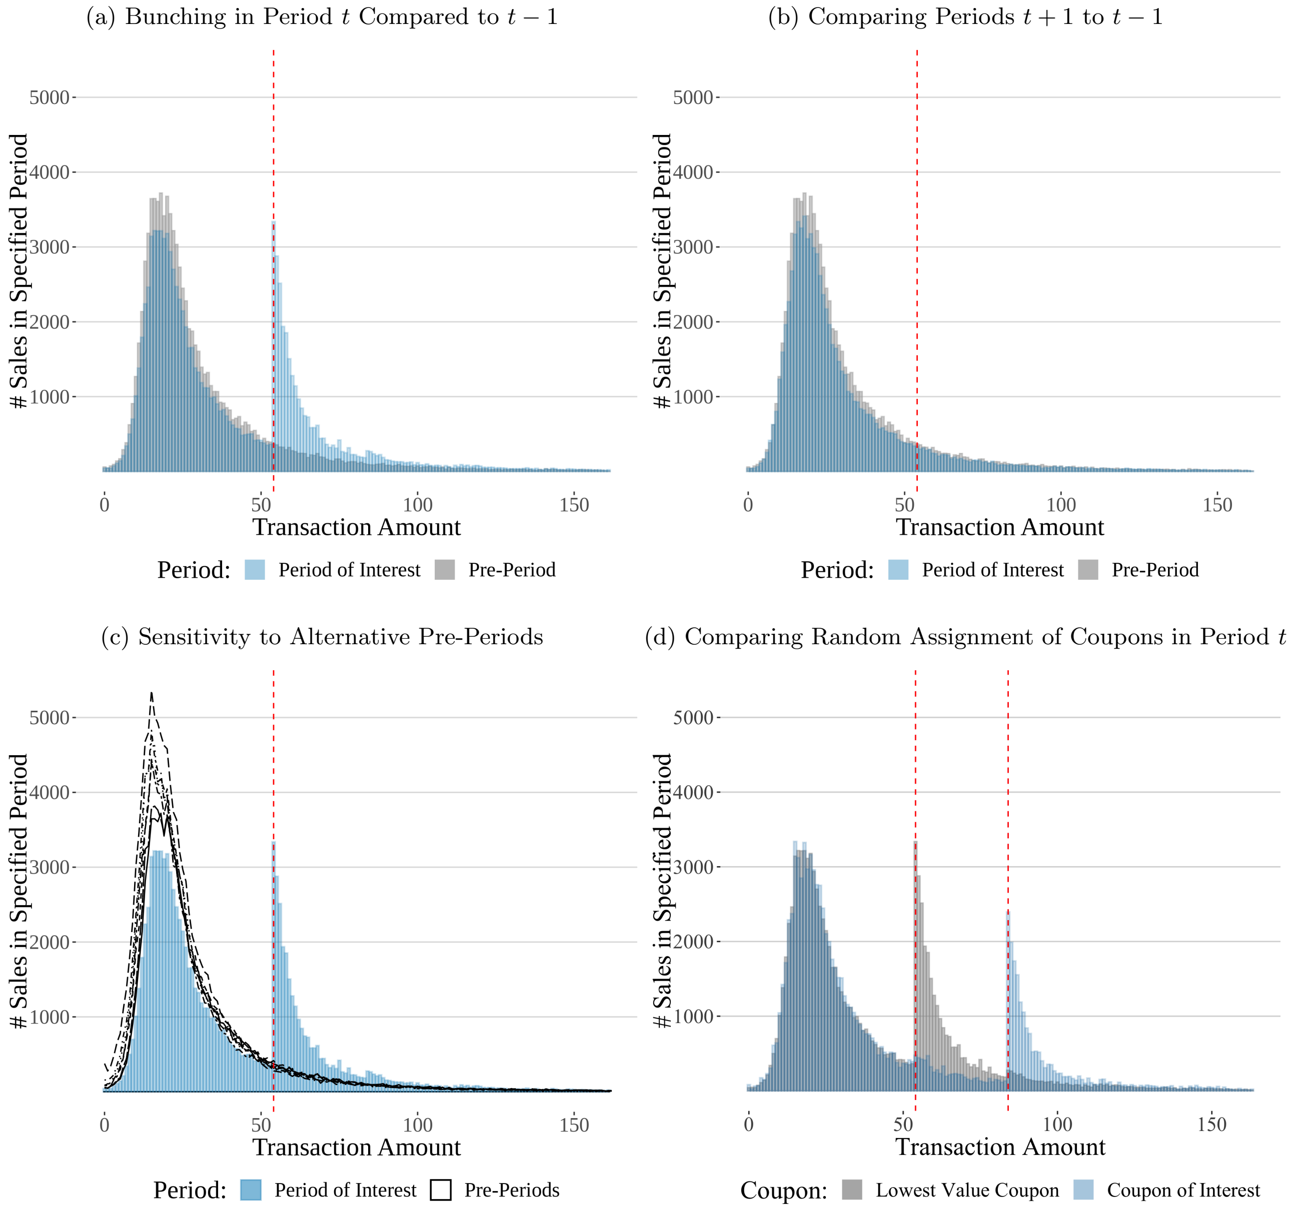

Notes: This figure illustrates the bunching estimator by comparing the distribution of food delivery spending between periods around the time the coupons were distributed. Panels (a)-(c) show the distribution of spending for the 54-18 coupon in City A. Panel (a) shows the distribution of spesnding during the coupon wave. Panel (b) shows the distribution of spending in the period immediately after coupons were distributed. In both (a) and (b) the pre-period $t-1$ distribution is shown for reference. Panel (c) illustrates the sensitivity of the bunching to different pre-periods by comparing the distribution in the coupon period to seven different pre-periods ($t-1$ through $t-7$). The analogous figures covering all of the spending categories covered by the coupon is available in the Appendix (see Figure OA.17). Lastly, panel (d) compares the distribution of spending during the coupon wave between consumers who were randomly assigned different coupons (either the 54-18 or the 84-28 coupon) in Wave 2 in City A; there is clear bunching at the coupon thresholds for each group of consumers, and there is greater overall spending for the consumers randomly assigned to the higher-threshold/higher-discount coupon.

Figure 2
Graphical Model

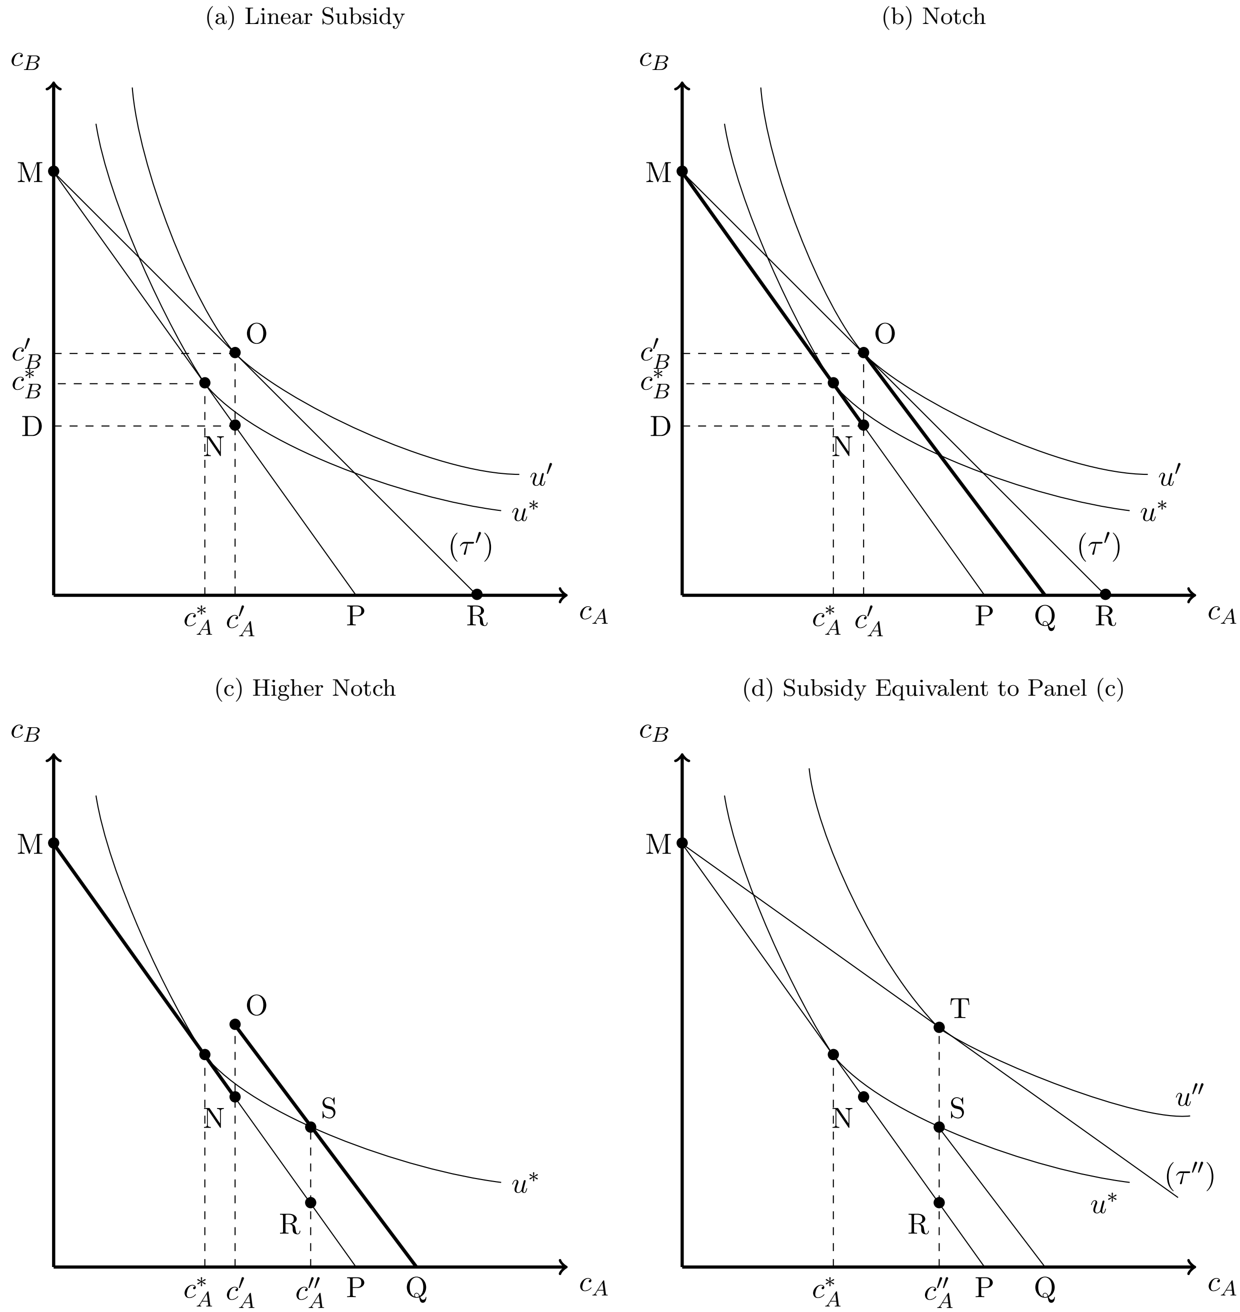

Notes: This figure presents a simple two-good graphical model to reassess the economics of notches versus linear subsidies. In Panel (a), the consumer responds to a linear subsidy that reduces the price of good A by a factor $(1-\tau')$. This rotates the budget constraint and leads to new choices $c_A^{'}$ 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 3

Assessing Intertemporal Substitution Using MPC^{coupon} Estimates Over Time

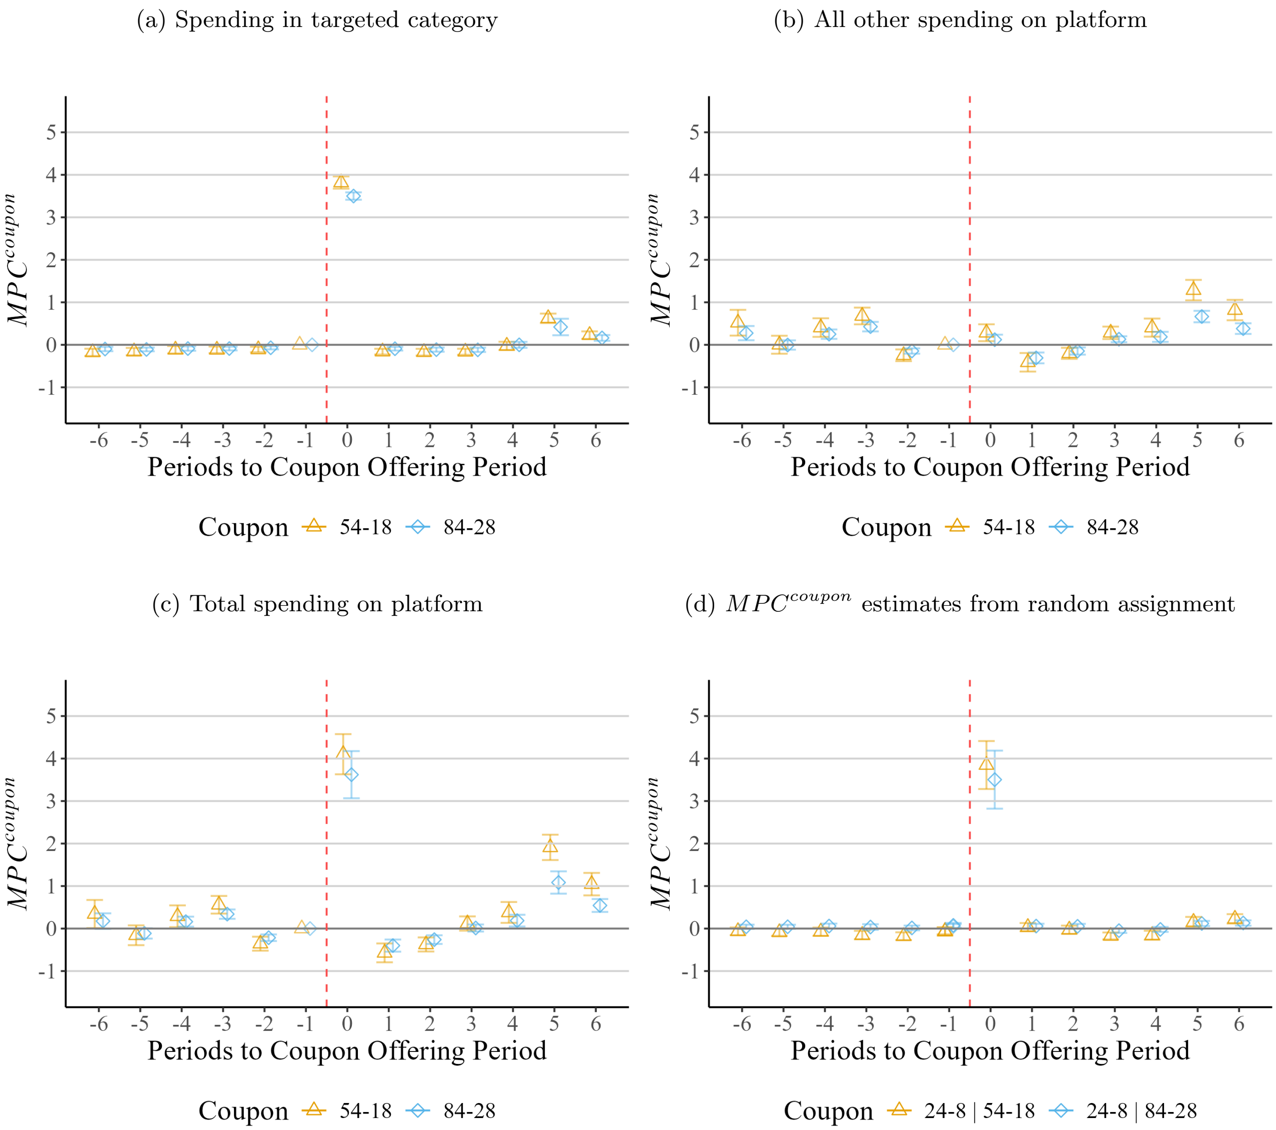

Notes: This figure reports MPC^{coupon} estimates over time for the Supermarket coupons distributed in wave 2 in City A. Panel (a) reports results for spending in the targeted spending category. Panel (b) reports results for all other spending on the platform, and Panel (c) reports results for total spending on the platform. Panel (d) reports results comparing consumers randomly assigned different coupons, following the same procedure used in Table 2 for all of the time periods. The confidence intervals are built with bootstrapped standard errors based on 1000 replications of a cluster-based bootstrap procedure that resamples the ¥1 bins of transactions with replacement.