

Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore[†]

By SUMIT AGARWAL AND WENLAN QIAN*

This paper uses a unique panel dataset of consumer financial transactions to study how consumers respond to an exogenous unanticipated income shock. Consumption rose significantly after the fiscal policy announcement: during the ten subsequent months, for each \$1 received, consumers on average spent \$0.80. We find a strong announcement effect—19 percent of the response occurs during the first two-month announcement period via credit cards. Subsequently, consumers switched to debit cards after disbursement before finally increasing spending on credit cards in the later months. Consumers with low liquid assets or with low credit card limit experienced stronger consumption responses. (JEL D12, D14, E21)

Jappelli and Pistaferri (2010) develop a theoretical framework that has several predictions for consumption response to unanticipated and anticipated income shocks depending on the persistence of the shocks and the degree of completeness of credit and insurance markets. Specifically, they argue that while consumption should not respond to anticipated income changes, it should respond to unanticipated income changes. While the literature on anticipated shocks is large, very few papers study unanticipated shocks, mainly due to the difficulty in identifying income shocks that are genuinely exogenous and unanticipated.¹

* Agarwal: NUS Business School, National University of Singapore, 15 Kent Ridge Drive, Singapore 119245, Singapore (e-mail: ushakri@yahoo.com); Qian: NUS Business School, National University of Singapore, 15 Kent Ridge Drive, Singapore 119245, Singapore (e-mail: wenlan.qian@nus.edu.sg). For helpful comments, we thank two anonymous referees, Chris Carroll, Souphala Chomsisengphet, Steve Davis, Takeo Hoshi, Tullio Jappelli, David Laibson, Alexander Ljungqvist, Brigitte Madrian, Neale Mahoney, Jessica Pan, Jonathan Parker, Ivan Png, Nagpurnanand Prabhala, Tarun Ramadorai, David Reeb, Amit Seru, Nick Souleles, Bernard Yeung, and seminar participants at AEA annual meeting, European Conference on Household Finance, 7th Singapore International Conference on Finance, Econometric Society Asian Meeting, Asian Bureau of Finance and Economic Research Inaugural Meeting, the World Bank, Oxford University, Johns Hopkins University, Nanyang Technological University, and National University of Singapore. We gratefully acknowledge funding support from the NBER Household Finance Working Group through the Alfred P. Sloan Foundation. Qian also thanks National University of Singapore for research funding (grant: R-315-000-100-112). The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <http://dx.doi.org/10.1257/aer.104.12.4205> to visit the article page for additional materials and author disclosure statement(s).

¹Some recent papers that study the consumption response to anticipated and temporary changes in income include Carroll (1992, 1997); Parker (1999); Shapiro and Slemrod (1995, 2003a, 2003b); Souleles (1999, 2000, 2002); Hsieh (2003); Stephens (2003, 2006, 2008); Johnson, Parker, and Souleles (2006); Agarwal, Liu, and Souleles (2007); Stephens and Unayama (2011); and Parker et al. (2013). The literature has explained the response to consumption of these expected income shocks through models of liquidity constraints and precautionary savings. A few papers that study the consumption response to unanticipated and temporary changes in income are Woplin (1982); Paxson (1993); Gruber (1997); and Jappelli and Pistaferri (2014). For a review of the literature, see Browning and Crossley (2001a) and Jappelli and Pistaferri (2010).

In this paper, we study a unique policy experiment by the Singapore government that is exogenous and allows one to distinguish (or to estimate) an announcement effect and a disbursement effect. On February 18, 2011, in a surprise announcement as part of the budget speech, the government announced the Growth Dividend Program. It constituted a one-time cash payout of US\$1.17 billion, ranging from US\$78 to US\$702 per person, to 2.5 million adult Singaporeans. The package was 0.5 percent of the annual gross domestic product (GDP) of Singapore in 2011, and was equivalent to 12 percent of Singapore's monthly aggregate household consumption expenditure in 2011. The money was distributed at the end of April, by which time the exact payout was known and expected.

We use a unique panel dataset of consumer financial transactions to study how consumers respond to this exogenous, unanticipated income shock. Specifically, we use a representative sample of more than 180,000 consumers in Singapore and study how their credit card, debit card, and bank checking account spending behavior responded to the positive income shock.² Close to 30 percent of aggregate personal consumption in the country is purchased using credit and debit cards.³ The richness of our data gives us the opportunity to study the response in credit card spending, debit card spending, the change in credit card debt, as well as the change in banking transaction behavior in the ten months following the policy announcement. Our analysis is based on a difference-in-differences identification that exploits the program's qualification criteria—foreigners did not qualify for the program and thus comprise the control group in the study. Unlike in other countries, foreigners in Singapore constitute close to 40 percent of the population, and are well represented across age, income, wealth, and other demographics.

We estimate a distributed lag model using the announcement date of the Growth Dividend Program as the exogenous event and obtain the impulse response of credit card spending, debit card spending, and credit card debt. Our findings are summarized as follows. First, recipients' consumption rose significantly after the fiscal policy announcement: for each \$1 received, consumers on average spent \$0.80 (aggregated across different financial accounts) during the ten months after the announcement. Prior to the announcement, on the other hand, there is no difference in the consumption trend between the treatment group and the control group. Second, we find a strong announcement effect: consumers started to increase spending during the two-month announcement period before the cash payout. For each \$1 received, \$0.15 (or equivalently 19 percent of the total consumption response) were consumed during the announcement period. Third, the consumption response

²As advocated by Gross and Souleles (2002) in the case of the United States, consumer credit also plays an important role in Singapore. More than one-third of consumers in the country have a credit card, and the total credit card debt as a percentage of GDP was over 2 percent in Singapore in 2011 and virtually everybody in Singapore has a debit card (Department of Statistics Singapore 2012).

³The remaining 70 percent of consumption is transacted via checks, direct transfers, and cash. Consumers with recurring payments such as mortgage payments, rent payments, and auto loan payments use instruments such as checks and direct deposit. We confirm this using our credit and debit transaction level data. Looking through the transaction category codes, merchant names, transaction types, we do not find a *single* transaction for mortgage, rent, and auto loan payments in over 18 million debit card and credit card transactions. Hence, we conclude that these recurring payments use checks and direct deposits. We empirically also confirm that these recurring payments conducted using instruments including checks and direct deposits are not sensitive to transitory income changes. Jappelli, Pischke, and Souleles (1998) found that people with bank cards were better able to smooth their consumption to income fluctuations than were people without bank cards.

was concentrated in debit card (25 percent of the total response) and credit card (75 percent of the total response) spending. We find that consumers started spending via credit cards during the announcement period, then switched to debit cards after disbursement, before finally increasing their credit card use significantly. Consistent with Agarwal, Liu, and Souleles (2007), there is a moderate decrease in credit card debt. Lastly, consumption response was heterogeneous across spending categories and across individuals. Consumption rose primarily in the discretionary category or small durable goods. Constrained consumers, measured as those with low liquid assets or with low credit card limit, showed stronger consumption responses.

We conduct a series of robustness tests. First, to address the concern that foreigners differ from Singaporeans in unobservable ways that may affect their spending behavior, we (i) restrict the control group to foreigners who are more similar to Singaporeans in ethnic and cultural backgrounds, and (ii) completely drop the foreigners from the sample and perform tests by exploiting the heterogeneity in the payout amount within the treatment group. Second, we restrict our treatment group to those unaffected by other potential treatments at the same time (e.g., other government programs and the annual bonus payout). Lastly, we investigate the robustness of our statistical inference—consistency of standard errors—and conduct our tests using alternative specifications. The results from these robustness tests are qualitatively and quantitatively similar to those in the main analysis.

This paper contributes to the literature in the following ways. As Jappelli and Pistaferri (2010) document, anticipated and unanticipated income shocks bear different implications for the consumption response. Permanent income hypothesis (PIH) suggests that consumption should respond to an unexpected increase in income: the magnitude of the consumption response is equal to the real interest rate in a complete market with an infinite horizon but will be higher when the horizon is finite. Moreover, the consumption response could be significant when consumers face borrowing constraints (Zeldes 1989a) or when precautionary saving motives are strong as the unanticipated income increase reduces the income uncertainty and encourages immediate spending (Zeldes 1989b; Carroll 1992, 1997). Prior studies (e.g., Poterba 1988) were unsuccessful in identifying the announcement effect primarily due to data weakness (e.g., low frequency or lack of cross-sectional variation) and the fact that a typical policy announcement was not a surprise. Other studies that use temporary job loss or illness as identification (e.g., Gruber 1997; Browning and Crossley 2001b; Gertler and Gruber 2002) are potentially subject to endogeneity and external validity concerns. Because the stimulus program we study has a surprise announcement date, we are able to show that consumption responds to an unanticipated income increase (during the announcement period). Since the income shock analyzed in this paper targets the general population (as opposed to specific groups often studied in the previous literature), it provides new, generalizable evidence on consumption response to unanticipated income shocks. This also has a significant economic implication. In our study, \$0.15 out of every \$1 received were consumed during the announcement period.⁴

⁴In a concurrent study, Jappelli and Padula (2014) study an unanticipated income shock in Italy and find consistent evidence of a significant consumption response.

Our paper is also closely related to studies of consumption response using micro-level data (e.g., Agarwal, Liu, and Souleles 2007). Both our study and theirs look at the dynamics of consumption and debt as a function of an income shock. Consistent with the existing findings, we find a significant spending response working through consumers' balance sheet, and we confirm that consumers with low liquid assets or low credit access respond more. Moreover, we document the dynamics of the consumption response across different spending instruments—a rise in credit card spending following the policy announcement, then the switch to (cheaper) debit card spending after disbursement of the stimulus, and finally the switch back to credit card use in the later months. This newly documented consumption response mechanism highlights that financial incentives drive (constrained) consumers' spending behavior. In addition, our results on the decomposition of the consumption response into different spending instruments imply that prior work based on micro-data from one single payment instrument (e.g., credit card) likely underestimates the spending response (to anticipated income shocks) due to limitations in their data.⁵ Specifically, they are unable to measure spending via debit cards, which we show accounts for a significant portion of the marginal propensity to consume (MPC), especially in the initial months after disbursement.

Finally, we add to the existing literature on the role of (financial) constraints in understanding consumption response to income shocks (Gross and Souleles 2002; Agarwal, Liu, and Souleles 2007; Leth-Petersen 2010). Our data allow us to identify several proxies of constraints such as the level of liquid assets and credit access, and our findings provide additional credibility to the result that constraints—whether labeled credit or liquidity constraints—are important for consumption response.

The rest of the paper flows as follows. Sections I and II discuss the fiscal policy experiment in Singapore and the data/econometric methodology respectively. The results appear in Sections III and IV, and Section V concludes.

I. The Growth Dividend Program in Singapore

The Ministry of Finance in Singapore announced on February 18, 2011 during the annual budget speech that in an attempt to share the nation's economic growth in 2010, the government would distribute a one-time payout of growth dividends to all Singaporeans over 21 years old in 2011. While the amount each Singaporean received depended on his or her wealth, a typical qualified Singaporean received between US\$428 and US\$624 in growth dividends. This payment represented a significant income bonus, corresponding to about 18 percent of monthly median income in Singapore in 2011. The program's payments totaled US\$1.17 billion, comparable to the size of the 2001 and 2008 US tax rebate.

Eligible Singaporeans received the payment by the end of April 2011, typically via direct bank transfer. The amount of the growth dividend an individual received was jointly determined by income and annual home value. The annual home value is the estimated annual rental revenues if the property were to be rented out, excluding the furniture, furnishings, and maintenance fees, and is determined by IRAS

⁵Recent work (Baker 2014; Gelman et al. 2014) starts to use richer administrative data beyond credit card datasets to uncover consumption response to income shocks.

(Singapore's tax authority) annually. We do not have data on the exact annual home value for each individual in our dataset, but we take advantage of the fact that the government uses the annual value of home criteria to identify less well-off Singaporeans living in government housing (known as HDB). Thus, we use the property type (HDB or private) together with income to identify the size of the growth dividend for each qualified Singaporean. In addition, adult Singaporean men who were serving or who had served in the army received an additional growth dividend of \$100 in recognition of their contribution to the nation. The average growth dividend amount that each qualified individual in our sample received was SG\$522 (US\$407). See Table A1 (in the online Appendix) for the exact payout schedule and how we proxy for the amount of the dividend received.⁶

Unlike other stimulus programs such as tax rebates in the United States, the Growth Dividend Program was unanticipated. We perform a thorough search of the newspaper articles related to the program, and find no discussion during the six-month period before the budget speech announcement. Within one week after the announcement, on the other hand, all major newspapers in Singapore had coverage of the stimulus program, highlighting the unanticipated nature of the program. The information became more salient in April shortly before disbursement, when the government (i) sent a letter to each qualified citizen concerning the exact amount of the cash payment, and (ii) provided an online calculator as well as a telephone hotline for questions on (the amount of) the cash payout program. This unique policy experiment allows us to distinguish between the announcement effect, which was a shock to the consumers and the disbursement effect, which was expected. Theory has different predictions for the two effects.

Other stimulus programs were announced at the same time in February 2011, but we focus on the growth dividend for several reasons. The Growth Dividend Program was significantly larger than the other stimulus packages. The program was unanticipated by the population. It also has features that allow better identification of the consumption response: it is the only one with cash payment (as opposed to, for example, an increase in the illiquid retirement account) that has a broad target population (i.e., all adult Singaporeans). This allows us to study the consumption response of the overall population by exploiting the untargeted foreigners to identify the counterfactuals. To control for the confounding effects of other stimulus packages, we drop from our analysis individuals who qualify for another cash stimulus package, the Workfare Special Bonus.⁷ We also perform other robustness checks to verify our results.

⁶ Admittedly, the approximation introduces measurement error in the amount of the cash payment the treatment group receives. This implies that our consumption and debt response estimates are downward biased. As a further robustness check, we also study the response by Singaporeans living in the private housing market for whom the benefit amount is observed with more accuracy and our results are qualitatively the same.

⁷ The qualification criteria for the Workfare Special Bonus were as follows: Singaporean citizens who were at least 35 years old by December 2010 and who worked for at least three months out of any six-month period in 2010 with a monthly income lower than SG\$1,700.

II. Data and Methodology

A. Data

We use in our analysis a unique, proprietary dataset obtained from the leading bank in Singapore that has more than four million customers, or 80 percent of the entire population in Singapore. As the largest bank in Singapore, it has more than twice the number of branches and over four times the number of automatic teller machines (ATMs) than the other major banks in Singapore. On the other hand, the typical banking fees and other costs are quite similar between our bank and the other major banks in Singapore (please refer to the online Appendix Table A2 for a detailed comparison on the banking facilities and fees among banks in Singapore). Although we do not have information on whether consumers have other banking relationships, our bank is likely the dominant bank for our sample consumers' daily financial needs due to its greater convenience and comparable banking fees.

Our sample contains consumer financial transactions data of more than 180,000 individuals, which is a random, representative sample of the bank's customers, in a 24-month period between 2010:04 and 2012:03. For each individual in our sample period, we have monthly statement information about each of their credit cards, debit cards, and checking accounts with the bank, including balance, total debit and credit amount (for checking accounts), spending (for credit and debit cards), credit card limit, credit card payment, and debt.⁸ At the disaggregated level, the data contain transaction-level information about each individual's credit card and debit card spending, including the transaction amount, transaction date, merchant name, and merchant category. The dataset also contains a rich set of demographics information about each individual, including age, gender, income, property type (HDB or private), property address zip code, nationality, ethnicity, and occupation.⁹

This dataset offers several advantages. Relative to traditional household spending datasets in the United States such as the Survey of Consumer Finance, our sample is larger with little measurement error, and it allows high frequency analysis. Compared to studies that use micro-level credit card data (e.g., Gross and Souleles 2002; Agarwal, Liu, and Souleles 2007; Aaronson, Agarwal, and French 2012), we have more complete information on the consumption of each individual in our sample.¹⁰ Rather than observing a single credit card account, we have information on every credit card, debit card, and checking account that individuals in the sample hold with the bank. Although we do not have information about financial instruments

⁸The specific banking products that we study (credit card, debit card, and bank checking account) are similar to those used in the United States. Consumers are typically eligible for obtaining a bank checking account, and they can conduct banking transactions using branches, Automatic Teller Machines (ATMs) (for cash withdrawals, transfers, or bill payment), checks, or online methods. The banking fees and other costs are quite standard as for a typical US bank, and are moreover comparable with banking costs at other major banks in Singapore. Debit cards are linked to the bank checking account, and debit card transactions are drawn on the bank account balance. Similarly, credit cards are granted upon application to consumers who have met the bank's criteria (e.g., income, age, and credit profile). One interesting difference for credit cards is that all credit card holders with the bank have the same prevailing interest rate of 24 percent per annum, regardless of the credit card limit.

⁹Unlike in the United States, where a zip code represents a wide area with a large population, a zip code in Singapore represents a building. A unique zip code is assigned to a single family house or for, say, a building with ten apartment units.

¹⁰Two recent studies have similar data to ours for the United States (Baker 2014 and Gelman et al. 2014).

individuals have with other banks in Singapore, we suspect the measurement error is negligible given the market share and representativeness of the bank. Furthermore, an average Singaporean consumer has three credit cards, which is also the number of credit cards an average consumer has in our dataset. In other words, we believe we are picking up the entire consumption of these consumers through their spending accounts with this bank. In addition, the richness of the transaction-level information as well as the individual demographics allows us to better understand heterogeneity in the consumption response to the positive income shock.

For our purpose, we aggregate the data at the individual-month level. Credit card spending is computed by adding monthly spending over all credit card accounts for each individual. Credit card debt is computed as the difference between the current month's credit card payment and the previous month's credit card balance. Debit card spending is computed by adding monthly spending over all debit card accounts for each individual. For the checking account, we compute the aggregate number of monthly debit (outflow) transactions for each individual. We exclude dormant/closed accounts and accounts that remained inactive (i.e., with no transactions) throughout the six months before the announcement of the Growth Dividend Program (i.e., 2010:08–2011:01). We use the first four months in our data (2010:04–2010:07) to identify consumers' pretreatment demographics, bank account, as well as consumer credit characteristics, to carry out the subsequent heterogeneity analysis. For a cleaner identification, we remove these four months from our sample. For a similar reason, we exclude months after 2011:11 to alleviate the potential concern of confounding effects due to seasonality and potential other government programs in the next year. As a result, the final sample period in our analysis is from 2010:08 to 2011:11.

Unlike the US stimulus policies under which clean identification stems from the random payout timing, in our policy experiment qualified consumers received the stimulus money at the same time. Instead, we use the difference-in-differences approach and rely on the untreated consumers—foreigners—to identify the consumption and debt response. This approach requires the control group to have the same spending and saving patterns as the treatment group in the pretreatment period so their behavior after the policy announcement constitutes a valid counterfactual.¹¹

Table 1 provides summary statistics of demographics and financial variables for the treatment and control groups in our sample. Panel A shows the demographics of the treatment and control groups. The control group (non-Singaporeans) is not directly comparable with the treatment group (Singaporeans) along several key dimensions. For example, the control group on average has a considerably higher income than the treatment group and is much less likely to live in government-subsidized housing (HDB). This suggests that the treatment group is less wealthy and may have a spending pattern inherently different from that of the control group.

¹¹ Foreigners in Singapore are an integral part of the population and are well represented across socioeconomic characteristics in the distribution. For example, according to the Population White Paper in Singapore, 27 percent of the foreigners in Singapore in 2011 are permanent residents, a subgroup already integrated into the population, who are on average young and hold a diploma of high education. The nonpermanent resident foreigners are a diverse group: 66 percent are on long-term working visas (20 percent are mid-level or high skilled workers and 46 percent are lower skilled workers), 14 percent are foreign domestic workers, and 21 percent are international students and family members of work visa holders.

TABLE 1—SUMMARY STATISTICS

	Treatment group		Control group		Diff. (5)
	Mean (1)	SD (2)	Mean (3)	SD (4)	
<i>Panel A. Demographics of the treatment and control groups</i>					
Age	44.09	10.57	40.31	8.45	-3.78***
Monthly income in 2010	6,053	8,861	7,795	11,395	1,742***
Female	0.42	0.49	0.30	0.46	-0.13***
Ethnicity					
Chinese	0.89	0.31	0.49	0.50	-0.40***
Malay	0.05	0.22	0.01	0.07	-0.04***
Indian	0.03	0.18	0.17	0.38	0.14***
Married	0.47	0.50	0.49	0.50	0.02***
Property type = HDB	0.70	0.46	0.55	0.50	-0.14***
\$D	522	213	0	0	
Observations	82,533		23,268		
	Matched treatment group		Matched control group		Diff.
Age	40.37	8.87	39.57	8.00	-0.79***
Monthly income in 2010	6,644	9,618	6,684	9,893	40
Female	0.38	0.49	0.35	0.48	-0.04***
Ethnicity					
Chinese	0.88	0.33	0.76	0.43	-0.12***
Malay	0.003	0.06	0.003	0.06	0.00
Indian	0.07	0.26	0.13	0.34	0.06***
Married	0.45	0.50	0.45	0.50	-0.00
Property type = HDB	0.66	0.47	0.65	0.48	-0.01
\$D	511	214	0	0	
Observations	36,989		10,567		

(Continued)

Furthermore, the amount of the growth dividend depends on the wealth level; thus, to reliably identify the policy effect, individuals in the treatment and control groups should also have comparable levels of wealth.

To this end, we construct a matched sample of Singaporeans (treatment) and foreigners (control) that are observationally similar. Specifically, we compute propensity scores based on a logistic regression using a rich set of income, wealth, as well as demographics information including age, gender, ethnicity, property type, and occupation (see Table A3 in the online Appendix for the logistic regression result). We perform the nearest-neighbor matching based on the computed propensity scores. After matching, the differences between the treatment and control groups in income and property type become statistically and economically indistinguishable from zero (panel A of Table 1). Differences in other characteristics also shrink significantly. In addition to the mean statistics, distributions of monthly income in 2010, age, and checking account balance of the treatment and control groups after matching are also similar and comparable (panel A of Figure 1). Therefore, we have a panel of reasonably balanced treatment and control individuals, which allows us to identify the average response as well as the dynamics of the treatment effect using the difference-in-differences approach.

Admittedly, the matched sample approach may not eliminate the unobservable differences between the Singaporeans and foreigners, which could affect their

TABLE 1—SUMMARY STATISTICS (*Continued*)

	Treatment group		Control group		Matched treatment group		Matched control group	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel B. Financial account information of the treatment and control groups</i>								
Credit card								
Spending	721	1,183	1,047	1,644	763	1,229	824	1,343
Cycle payment	701	17,193	1,043	2,485	736	6,258	798	1,843
Debt	719	1,958	843	2,268	667	1,921	655	1,945
Change in debt	6	490	7	652	6	514	4	551
Debit card								
Spending	216	507	235	542	225	519	190	476
Bank checking account								
No. debit transactions	15	12	16	11	15	12	12	10
No. ATM debit transactions	0.68	3.08	0.63	2.58	0.55	2.68	0.44	2.19
No. branch debit transactions	0.32	0.88	0.27	0.74	0.28	0.81	0.24	0.71
No. online debit transactions	0.26	0.71	0.29	0.72	0.32	0.78	0.30	0.73
Month-end balance	16,036	21,823	14,320	20,531	15,513	21,391	14,083	20,076
Total (card) spending	937	1,290	1,282	1,770	988	1,341	1,014	1,443
Total spending on supermarket	56	119	95	181	59	122	67	144
Total spending on service	251	461	287	541	261	473	257	486
Total spending on dining	66	194	113	292	74	206	79	230
Total spending on entertainment	49	167	52	147	48	157	39	127
Total spending on apparel	91	270	141	349	97	280	107	296
Total spending on travel	140	356	200	526	145	365	148	420
Total spending on small durable goods	90	277	68	236	73	246	75	247
Total spending online	18	79	23	92	19	82	19	84
Observations	1,893,217		512,213		845,339		233,197	

Notes: This table reports the summary statistics of our treatment and control sample, both before and after propensity score matching (based on the nearest neighbor). The treatment sample consists of individuals who qualify for the Growth Dividend Program (but not for other cash stimulus packages such as the Workfare Special Bonus), and the control sample consists of all non-Singaporeans. We also exclude individuals/accounts that are dormant or closed or had no transaction activity during the six-month period before the policy announcement. Panel A shows the comparison of demographics between the treatment and control groups, both before and after propensity score matching. Panel B shows the comparison of credit card and debit card spending, credit card debt, and bank checking account balance information between the treatment and control groups in the 24-month sample period (2010:04–2012:03), both before and after propensity score matching. *SD* is the Growth Dividend amount individuals received in the treatment group. Credit card spending is computed by adding monthly spending over all credit card accounts for each individual. Credit card debt is computed as the difference between the current month's credit card payment and the previous month's credit card balance. Credit card cycle payment is the payment to the most recent credit card statement. Debit card spending is computed by adding monthly spending over all debit card accounts for each individual. For the checking account, we compute No. debit transactions as the aggregate number of debit (outflow) transactions for each individual every month. No. ATM debit transactions/No. branch debit transactions/No. online debit transactions are the number of debit transactions at Automated Teller Machines (ATMs), in branches, or via online transactions, respectively. Total card spending is the sum of debit card spending and credit card spending for each individual in a month. All the dollar amounts are in the local currency (SG\$), and SG\$1 = US\$0.78 as of February 2011.

consumption patterns. In our analysis, we will explicitly test for any difference between the treatment and control groups in the consumption and debt trends during the pretreatment period. We further carry out various robustness checks to validate our matched sample approach. In addition, we verify the external validity of our results by carrying out the difference-in-differences regressions in the full (unmatched) sample.

We first plot, in panel B of Figure 1, the unconditional mean total card spending of the treatment and control group in the matched sample during the period of 2010:04–2012:03. On average, the treatment group has higher total spending than the control group. Moreover, the difference in total spending between the treatment

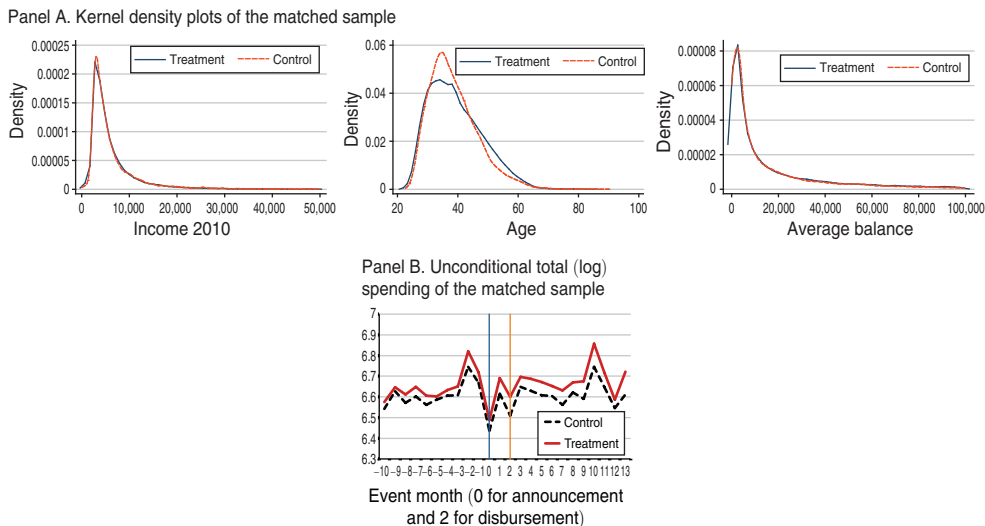


FIGURE 1. SUMMARY STATISTICS OF THE MATCHED SAMPLE

Notes: Panel A shows the treatment group and the control group comparison of distributions of average monthly income in 2010, age, and average bank checking account balance during the period 2010:04–2010:07, after the propensity score matching. Panel B shows (logarithm of) the unconditional mean of total spending of the treatment and control group during the period from 2010:04 to 2012:03.

group and the control group before the announcement of the Growth Dividend Program remains constant, which confirms the underlying identifying assumption of a parallel trend. Note that the gap between the treatment group and the control group visibly increases after the program, which provides the first suggestive evidence of the consumer spending response to the income shock.

B. Methodology

We analyze the response in spending and debt using a difference-in-differences regression methodology. The treatment group corresponds to Singaporeans who are entitled to the growth dividend payout, and the control group corresponds to foreigners in the country. The pretreatment period is from 2010:08 to 2011:01 (six months), and the post-treatment period is from 2011:02 to 2011:11 (ten months).

First, we study the average monthly response to the stimulus using the following specification:

$$(1) \quad Y_{i,t} = \beta_{pre} \times \$D_i \times 1_{pre} + \beta_{post} \times \$D_i \times 1_{post} + \alpha_i + \delta_t + \epsilon_{i,t}.$$

Following Agarwal, Liu, and Souleles (2007) and Aaronson, Agarwal, and French (2012), the dependent variable $Y_{i,t}$ is the dollar amount of total card spending (further decomposed into debit card spending and credit card spending), or change in credit card debt for individual i at the end of month t . $\$D_i$ is the amount of the growth dividend that individual i in the treatment group received, and is equal to 0 for the control group. 1_{pre} is a binary variable equal to 1 for the four months before

the announcement of the Growth Dividend Program (i.e., 2010:10–2011:01).¹² 1_{post} is a binary variable equal to 1 for the months after the announcement of the Growth Dividend Program (i.e., $\geq 2011:02$). δ_t is the year-month dummy, used to absorb the seasonal variation in consumption expenditures as well as the average of all other concurrent aggregate factors. α_i is the individual dummy included to absorb differences in consumption preferences at the individual level. Standard errors in all regression analyses are clustered at the individual level.

β_{post} in equation (1) captures the average monthly post-policy spending (or debt change) response per dollar received for a treated individual (compared to the benchmark period, i.e., from 2010:08 to 2010:09), relative to the post-policy change in spending (or debt change) of the control group. On the other hand, β_{pre} measures the difference in the spending (or debt change) trend between the treatment group and the control group during the four pretreatment months (compared to the benchmark period). Validity of our difference-in-differences research design requires β_{pre} to be statistically and economically indistinguishable from zero.

We also divide the post-policy window into the announcement period and the disbursement period, to compare the policy effect in these two windows separately.

$$(2) \quad Y_{i,t} = \beta_{pre} \times \$D_i \times 1_{pre} + \beta_a \times \$D_i \times 1_{announce} + \beta_d \times \$D_i \times 1_{disburse} + \alpha_i + \delta_t + \epsilon_{i,t}.$$

Specifically, $1_{announce}$ is a binary variable equal to 1 for the two months during the announcement window (2011:02–2011:03), and $1_{disburse}$ is a binary variable equal to 1 for the months after the disbursement of the growth dividend (i.e., $\geq 2011:04$). Therefore, the coefficients β_a and β_d in equation (2) capture the average monthly spending (or debt change) response per dollar received, relative to the change in spending (or debt change) of the control group, for a treated individual during the announcement period and the disbursement period, respectively.

In addition, we study the dynamics of the spending (or debt change) response. Specifically, we estimate the following distributed lag model:

$$(3) \quad Y_{i,t} = \sum_{s=-4}^9 \beta_s \times \$D_i \times 1_{month\ s} + \alpha_i + \delta_t + \epsilon_{i,t}.$$

Following Agarwal, Liu, and Souleles (2007), the results can be interpreted as an event study. The coefficient β_0 measures the immediate dollar response in spending (or debt change) per dollar dividend expected during the announcement month. The *marginal* coefficients β_1, \dots, β_9 measure the *additional* marginal responses one month, ..., and nine months after the announcement, respectively. Similarly, coefficients $\beta_{-4}, \dots, \beta_{-1}$ capture the difference of trends in spending and debt change between the treatment group and the control group in each of the four pretreatment months.

¹²2010:08 and 2010:09 are absorbed to reliably identify the benchmark spending/debt pattern in our estimation. We have also tried the specification where we only absorb 2010:08 and add 2010:09 as another pretreatment month ($s = -5$) in the estimation and the results remain to hold.

To gauge the expansionary impact of the fiscal stimulus, we define the *cumulative* coefficients $b_s \equiv \sum_{t=-4}^s \beta_t$ that describe the cumulative response in spending (or debt change) after s months, $s \leq 9$. Note that the coefficient b_s captures the cumulative response of spending and debt change from month -4 (i.e., four months before announcement). Thus, the cumulative effect of the spending (or debt change) at month s upon announcement is $b_s - b_{-1} (\equiv \sum_0^s \beta_t)$, $s \geq 0$. For instance, if spending rises by $\beta_0 = \$0.06$ on \$1 of growth dividend in the announcement month and after one month spending rises by $\beta_1 = \$0.09$ on \$1 of growth dividend, then the cumulative spending effect after month 1 is $b_1 - b_{-1} = \$0.15$ on \$1 of growth dividend. The response of debt change is of independent interest and can also help shed light on the spending response. On the other hand, b_{-4}, \dots, b_{-1} measure the cumulative spending (or debt change) differences between the treatment group and the control group by month $-4, \dots, -1$ in the pretreatment period, and we expect them to be economically and statistically insignificant.

We also study the heterogeneity in the response to the growth dividend across different groups of individuals (e.g., constrained versus unconstrained consumers) using the following specification:

$$(4) \quad Y_{i,t} = \sum_{s=-4}^9 \beta_s \times \$D_i \times 1_{month\ s} + \sum_{s=-4}^9 \beta_{g1,s} \times 1_{g1} \times \$D_i \times 1_{month\ s} + \dots \\ + \sum_{s=-4}^9 \beta_{g(N-1),s} \times 1_{g(N-1)} \times \$D_i \times 1_{month\ s} + \alpha_i + \delta_t + \epsilon_{i,t},$$

where N is the number of subgroups of consumers that we decompose into ($g1$ stands for the first group, \dots , $g(N-1)$ stands for the $(N-1)$ th group, and the N th group is the absorbed group and thus unshown in equation (4)).

III. Main Results

We begin by estimating the average response of spending (in various financial accounts) and credit card debt change to the Growth Dividend Program. To sharpen the results, we later analyze dynamics using a distributed lag model and study response heterogeneity across different spending categories and across different types of individuals. In the main analysis, we focus on the matched sample in the period from six months before to ten months after the announcement of the Growth Dividend Program (2010:08–2011:11). To further address the possibility that individuals spend via financial instruments issued by other banks, we include in our matched sample analysis only individuals who have a bank account, debit card, and credit card account with the bank at the same time.

A. The Average Response of Debit Card and Credit Card Spending and Credit Card Debt

Panel A of Table 2 shows results on the average response by applying equation (1) to spending and credit card debt change. Since 1_{pre} is a binary variable equal to 1 for the four months before the announcement of the Growth Dividend Program (i.e., 2010:10–2011:01), the coefficients on $\$D_i \times 1_{pre}$ measure the difference in spending

TABLE 2—THE AVERAGE SPENDING AND DEBT RESPONSE TO THE STIMULUS PROGRAM

	Total card spending (1)	Debit card spending (2)	Credit card spending (3)	Credit card debt change (4)
<i>Panel A</i>				
$\$D \times 1_{pre}$	−0.004 (−0.19)	0.006 (0.60)	−0.010 (−0.56)	0.001 (0.06)
$\$D \times 1_{post}$	0.080*** (3.60)	0.026** (2.40)	0.053*** (2.84)	−0.010 (−1.02)
Constant	1,123.426*** (161.88)	490.350*** (144.24)	633.076*** (106.27)	16.774*** (4.25)
R^2	0.541	0.499	0.534	0.032
<i>Panel B</i>				
$\$D \times 1_{pre}$	−0.004 (−0.19)	0.006 (0.60)	−0.010 (−0.56)	0.001 (0.06)
$\$D \times 1_{announce}$	0.074*** (2.78)	0.013 (0.99)	0.061*** (2.72)	−0.009 (−0.70)
$\$D \times 1_{disburse}$	0.081*** (3.52)	0.030*** (2.61)	0.051*** (2.62)	−0.010 (−1.01)
Constant	1,123.426*** (161.88)	490.351*** (144.24)	633.075*** (106.27)	16.774*** (4.25)
Fixed effects		Individual, year-month		
R^2	0.541	0.499	0.534	0.032

Notes: This table shows the average spending and debt (change) response (equations (1) and (2)) of the matched sample in the period from 2010:08 to 2011:11. Panel A presents the estimation results of equation (1), and panel B shows the estimation results of equation (2). $\$D$ is the amount of the growth dividend received for the treatment group, and is equal to 0 for the control group. 1_{pre} is a binary variable equal to 1 for the four months before the announcement (i.e., 2010:10–2011:01). 1_{post} is a binary variable equal to 1 for the months after the announcement of the Growth Dividend Program (i.e., \geq 2011:02). $1_{announce}$ is a binary variable equal to 1 for the months during the announcement window (2011:02–2011:03), and $1_{disburse}$ is a binary variable equal to 1 for the months after the disbursement of the growth dividends (i.e., \geq 2011:04). Please refer to Table 1 for definitions of other variables. Individual and year-month fixed effects are included, and standard errors are clustered at the individual level. T -statistics are reported in parentheses under the coefficient estimates.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

(or credit card debt change) per dollar of the expected growth dividend amount, compared to the first two months in our sample period (2010:08–2010:09), between the treatment and control groups in those four pretreatment months. Similarly, the coefficients on $\$D_i \times 1_{post}$ capture the spending (or credit card debt change) response after the announcement compared to the first two months in our sample period.

The first column shows the average response of monthly total card spending (i.e., debit card spending + credit card spending) of the treatment group. Overall, individuals in the treatment group increased their card spending by \$0.08 per month for every \$1 of growth dividend received. The effect is both statistically and economically significant, and it corresponds to a total increase of \$0.80 per \$1 received in the ten-month period after announcement.¹³ About two-thirds of the total spending

¹³The cumulative confidence interval is wide in the later months so a conservative estimate of the cumulative spending effect is smaller at such horizons. Nevertheless, even the muted magnitude of the cumulative spending response is comparable to that found in most of the previous literature.

increase after the stimulus program announcement are attributable to the spending increase on credit cards (\$0.53 per \$1 received, column 3 of panel A of Table 2), and one-third is due to spending on debit cards (\$0.26 per \$1 received, column 2 of panel A of Table 2). The coefficient on the credit card debt change is -0.01 (column 4 of panel A), which suggests that credit card debt experienced a \$0.01 decrease per month or a \$0.10 decrease in total per \$1 received for the treatment group in the ten months after announcement. But the effect is statistically insignificant.

In all four columns in panel A of Table 2, coefficient estimates on the pretreatment period variable $SD_i \times 1_{pre}$ are both economically small and statistically insignificant. For example, for each growth dividend dollar expected, the treatment group's monthly total card spending is on average \$0.004 *less* than the control group in the four months before the program announcement and is statistically insignificant (p -value = 0.852). To interpret, these results suggest that before the stimulus program, there are no differences in spending and credit card debt change patterns between the matched Singaporeans and foreigners. This provides strong evidence in support of our research design: the matched sample of Singaporeans (treatment) and foreigners (control) is balanced and homogeneous (in their spending and debt change trend), and the differences in spending and credit card debt change after announcement, indeed measure the treatment group's response to the income shock.

B. Announcement versus Disbursement Effect

The Growth Dividend Program was a one-time stimulus program that was *unanticipated* by the population in Singapore: the program was announced in February 2011, two months before qualified Singaporeans received the payments in April 2011. As a result, we can investigate the announcement effect separately from the disbursement effect. The prior literature (e.g., Poterba 1988) has not been successful at estimating the announcement effect because the announcement was not a surprise. However, the life-cycle theory has a clear prediction that consumers should respond to the announcement (of an unanticipated income shock). Our setting is the first to cleanly test this theory and we thus estimate equation (2) by decomposing the post-policy window into the announcement period and the disbursement period (panel B of Table 2).

We find a significant increase in total card spending in both windows: individuals spent \$0.074 per month for every \$1 expected in the two-month announcement period and \$0.081 per month for every \$1 received during the disbursement period. We cannot reject the hypothesis that the monthly total card spending response is the same between the announcement and the disbursement period (p -value = 0.737). Interestingly, there is a significant difference in the means of spending for the two windows. For the announcement period, the increase in spending is primarily concentrated in credit cards (\$0.061 per month for \$1 expected), while there is no statistically or economically significant change in debit card spending for the treatment group during this period. Debit card spending increased mostly in the disbursement window, and the credit card spending continued to be high for the treatment group during the disbursement window. There is little difference in credit card debt change in the announcement and the disbursement periods, as both coefficients are statistically insignificant.

In summary, our results add two new findings to the literature on consumption response to stimulus programs/income shocks. First, consumers started spending the stimulus money upon announcement of the program (i.e., a period when the income increase is unanticipated). The announcement effect is significant: compared to the disbursement window, consumers increased monthly spending in the announcement period by a similar amount. This is consistent with life-cycle model predictions of consumption response to unanticipated income shocks.¹⁴ Moreover, prior literature that evaluates similar fiscal policies likely underestimates the consumption response since a typical stimulus program being studied does not have a well-defined announcement period. Second, consumption responded in the announcement period primarily through credit card use. Upon receiving the dividend, they used their debit cards along with credit cards to increase spending. These results also reveal the role of consumer credit in facilitating the consumption response in the announcement period to the unanticipated income shock (Agarwal, Liu, and Souleles 2007; Johnson, Parker, and Souleles 2006). Consumers largely “borrowed” from the future stimulus money and started spending immediately upon announcement.

C. Response in Bank Checking Accounts

We next study debit transactions in consumers’ bank checking accounts before and after the announcement of the Growth Dividend Program. Because we do not have transaction-level data on debit transactions in the checking accounts, we first use the number of debit transactions as the dependent variable to investigate whether consumers in the treatment group increased the number of debit transactions significantly after the program. Column 1 of Table 3 shows the regression results. There is no significant change in the number of checking account debit transactions after the stimulus program for the treatment group. We also perform analysis on the one specific debit transaction (ATM) that is informative on cash spending. There is no significant change in the ATM transaction activity for the treatment group after the program (column 2 of Table 3).

We perform an additional test by inferring the amount of monthly cash/check spending for each individual. By assuming that individuals deposit their monthly income into and pay their credit card balance from this particular bank’s account, we estimate (a noisy measure of) the monthly cash/check spending as bank balance at the start of the month + income – total card spending – bank balance at the end of the month. In column 3 of Table 3, we show that there is no difference in the change of cash/check spending between the treatment group and the control group, which is consistent with the finding on the number of bank transactions. Taken together, these results suggest that consumers most likely increased their spending through card spending, either with debit or credit cards.

¹⁴As a further test, we identify consumers who likely have precautionary saving motives who should immediately increase spending (especially in the presence of market incompleteness), as the unanticipated shock reduces the income uncertainty and their precautionary saving motive. Consistently, we show (in detail) in the online Appendix that those consumers have a strong announcement effect.

TABLE 3—SPENDING RESPONSE FROM THE BANK CHECKING ACCOUNT

	In no. of debit transactions (1)	In no. of ATM transactions (2)	Inferred cash/check spending (3)
$1_{treatment} \times 1_{pre}$	0.005 (1.14)	0.000 (0.04)	
$1_{treatment} \times 1_{post}$	0.005 (0.87)	-0.001 (-0.21)	
$\$D \times 1_{pre}$			0.270 (0.44)
$\$D \times 1_{post}$			0.196 (0.40)
Constant	2.788*** (1,381.65)	0.115*** (110.85)	4,601.547*** (32.39)
Fixed effects R^2	0.842	Individual, year-month 0.919	0.049

Notes: This table studies whether spending out of bank checking account changes for the treatment group after the announcement of the stimulus program in the matched sample during the period of 2010:08–2011:11. The dependent variables are (natural logarithms) of 1 plus the number of total debit transactions (column 1), (natural logarithms) of 1 plus the number of ATM transactions (column 2), (inferred) cash or check spending out of the bank checking account (column 3). By assuming that individuals deposit their monthly income into and pay their credit card balance from this particular bank's account, we estimate (a noisy measure of) the monthly cash/check spending as bank balance at the start of the month + income – total card spending – bank balance at the end of the month. 1_{pre} is a binary variable equal to 1 for the four months before the announcement (i.e., 2010:10–2011:01). 1_{post} is a binary variable equal to 1 for the months after the announcement of the Growth Dividend Program (i.e., $\geq 2011:02$). Individual and year-month fixed effects are included, and standard errors are clustered at the individual level. *T*-statistics are reported in parentheses under the coefficient estimates.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

To better understand the no-response in cash or check spending (which overall accounts for 70 percent of the entire consumption), we note that consumption using those instruments is primarily nondiscretionary. For example, people use cash or checks for big and recurring expenses such as tuition, mortgage, rent, and car loan payments (which they cannot pay using either debit or credit cards).¹⁵ As Singapore is one of the most expensive countries in the world to own a house or a car, a stimulus check of about \$500 is unlikely to trigger spending in big-ticket durables such as cars, which stands in contrast to the finding in studies on the US income shocks (e.g., Aaronson, Agarwal, and French 2012; and Parker et al. 2013).¹⁶

¹⁵We confirm this using our credit and debit transaction level data—looking through the transaction category codes, merchant names, and transaction types—and do not find a *single* transaction for mortgage, rent, and auto loan payments in over 18 million debit card and credit card transactions.

¹⁶Singapore is the most expensive country to own a car. To illustrate the cost of owning a car, consumers have to pay hefty taxes and license fees. Tax typically ranges two or three times the value of the car, and the average cost of having the license (to drive for ten years) is about SG\$60,000 in our post-treatment period. Adding together, the retail price of a *middle-class* car in Singapore is SG\$205,000 (US\$168,000) in 2011 (Ho, Png, and Reza 2014).

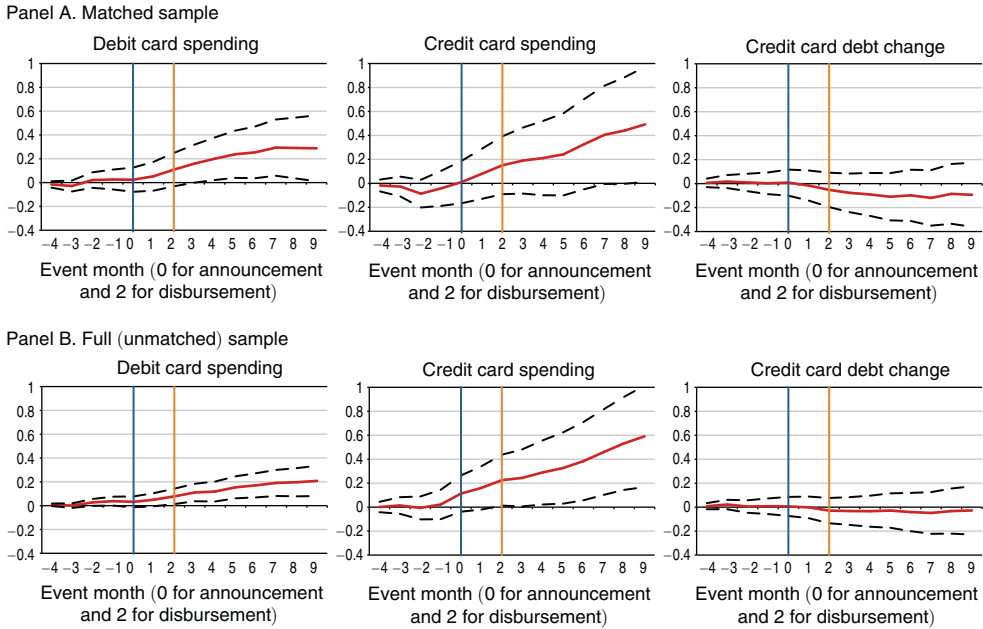


FIGURE 2. ESTIMATED SPENDING AND DEBT RESPONSE DYNAMICS

Notes: This figure plots the entire paths of cumulative coefficients b_s , $s = -4, -3, -2, \dots, 8, 9$, along with their corresponding 95 percent confidence intervals, of debit card and credit card spending as well as credit card debt change response as estimated from equation (3). The sample in panel A includes the matched treatment and control groups during the period of 2010:08–2011:11. Panel B reports results from the full unmatched sample in the same period (also does not require each consumer to have all three accounts—bank checking account, debit card, and credit card—with the bank). We estimate the full sample results with weighted least square regressions, using the propensity scores (from Table A3 in the online Appendix) as weights. The x -axis denotes the s th month after the announcement of the Growth Dividend Program, and the y -axis shows the dollar response (for every dollar received).

D. The Dynamics of the Spending and Debt Response

Results in Tables 2–3 show the average monthly response of spending and debt change to the stimulus program. In addition, to gauge the expansionary impact of the fiscal stimulus, we investigate the dynamic evolution of the spending and debt change response during the ten-month post-announcement period beginning from four months before the program announcement (equation (3)). Panel A of Figure 2 graphs the entire paths of cumulative coefficients b_s , $s = -4, -3, \dots, 8, 9$, and the dotted lines depict the corresponding 95 percent confidence intervals. Standard errors of the cumulative effects are calculated based on the standard errors of the marginal coefficients in the regression, which are clustered at the individual level. The results can be interpreted as an event study, with month 0 (2) being the time of announcement (disbursement). As noted before, the cumulative effect of the spending change at month s upon announcement is measured by $b_s - b_{-1}$.

Consistent with the regression results in Table 2, the cumulative spending and credit card debt change differences between the treatment group and the control group during the four-month pre-announcement period are insignificant both statistically and economically. Spending started to increase after the stimulus program announcement. By the end of nine months after the announcement month, the

cumulative increase in total card spending from the announcement month ($b_9 - b_{-1}$) is \$0.80 per \$1 received (p -value < 0.001).

Decomposing the total card spending into debit card and credit card spending gives more insight into the spending response dynamics. We find no spending response in debit card spending during the announcement period: the marginal effect for either of the two announcement months is insignificant for debit card spending. On the other hand, coefficients in the credit card spending regression are highly significant for both announcement months, suggesting that the total spending response in the announcement period is attributable to an increase in credit card spending ($b_1 - b_{-1} = 0.122$, p -value = 0.007). After disbursement, consumers in the treatment group primarily used their debit cards to increase their spending in the earlier period, since the marginal effect coefficients are statistically insignificant for credit card spending in months three and four after announcement. In the later period, the debit card spending increase gradually plateaued by month eight, but the credit card spending increase picked up again, and there is still evidence of a marginal increase in credit card spending in months eight and nine. A formal statistical test shows that debit card spending response is more front-loaded, as the cumulative increase in debit card spending in the first five months is larger than that in the last five months (p -value = 0.078). On the other hand, the cumulative increase in credit card spending in the first five months is statistically indistinguishable from that in the last five months (p -value = 0.714).

The point estimates on credit card debt change show that consumers started to reduce their debt when they received the growth dividend ($s = 2$) by \$0.036 per \$1 received, and the effect is statistically significant at the 5 percent level. After that, they continued to reduce their debt but the marginal effect coefficient estimates in the subsequent months become smaller and remain statistically insignificant. We also perform an additional statistical test: the difference between the cumulative credit card debt changes in the first five post-announcement months compared to that in the last five post-announcement months is $-\$0.09$ (per dollar received) and is statistically significant (p -value = 0.028). This suggests that the credit card debt decrease is short-lived and concentrated in the earlier period, especially after the disbursement of the growth dividend, before it stops in the second half of our post-policy period.

Taken together, the results in panel A of Figure 2 suggest that consumers in the treatment group responded strongly to the stimulus program upon announcement by increasing their spending via credit cards. We find a delayed spending response via debit cards, which occurred only after the payment of the stimulus money and then gradually plateaued over time. At the same time, consumers started to decrease their credit card debt as well as reducing their credit card spending in the early period after the disbursement. However, in the last few months of the ten-month treatment period, they stopped paying down their credit card debt and increased their credit card spending significantly again.

E. Full Sample Analysis

We perform the main analysis in the previous sections on a smaller sample in which the treatment group and control group are matched on several demographics variables. To ensure that the results can be generalized to the full sample, we repeat

the main analysis on the unmatched sample. In particular, we include all consumers who have at least one active account with the bank in the analysis. In the full sample, the treatment group (Singaporeans) and the control group (foreigners) are observationally different along key dimensions such as income and ethnicity (panel A of Table 1). To address the challenge in statistical inference related to an unbalanced sample, we exploit the estimated propensity scores (for the matched sample analysis) and include them as regression weights in the full sample difference-in-differences analysis. The rationale is to give a larger weight to those more similar foreigners in the control group (e.g., those with similar income, age, or cultural origins) in estimating the counterfactuals after the stimulus program. On the other hand, by using all treated Singaporeans in the analysis, we are able to speak to the external validity of our results in the matched sample analysis.¹⁷

Rather than reporting the regression coefficients, we plot the full dynamics of the spending and credit card debt change response. Panel B of Figure 2 shows the cumulative effects of debit card spending, credit card spending, as well as credit card debt change from four months before the stimulus announcement ($s = -4$). In the full sample, we observe a qualitatively and quantitatively similar effect in spending and credit card debt change. By month nine after the stimulus announcement, the total card spending has increased by \$0.80 per \$1 received (p -value < 0.001), out of which one-quarter is attributable to debit card spending. Credit card debt does not experience any significant change after announcement: the ten-month cumulative decrease in credit card debt is \$0.03 per \$1 received (p -value = 0.788). Moreover, consumption response exhibits the same pattern/dynamics as that estimated from the matched sample: consumers started spending during the announcement period mainly using credit cards, switched to debit cards upon disbursement, before finally increasing their credit card spending significantly. Similarly as before, we do not observe any differences in spending or credit card debt change between the treatment and control groups during the four pre-announcement months, which further validates our research design.

IV. Heterogeneity and Robustness

A. Heterogeneity of Spending and Debt Response across Consumers

We next study the dynamics of heterogeneous responses to the stimulus program across different consumers. Previous literature has shown that constrained consumers' consumption responds more strongly to positive income shocks (e.g., Agarwal, Liu, and Souleles 2007). We have a rich array of account-holder information including demographics and financial health information, which allows us to study the heterogeneous response of consumers in greater depth.¹⁸ Furthermore, our data allow us to understand differences in the full path of the consumers' spending responses

¹⁷The standard errors reported for the full sample analysis do not take into account the estimation error associated with estimating the propensity score, and the reported errors may thus be downward biased, making estimates appear more precise than they actually are.

¹⁸Earlier studies use demographics as proxies for liquidity constrained consumers. For instance, prior papers argue that young and old consumers are more likely to be liquidity constrained. Additionally, married consumers are less likely to be liquidity constrained.

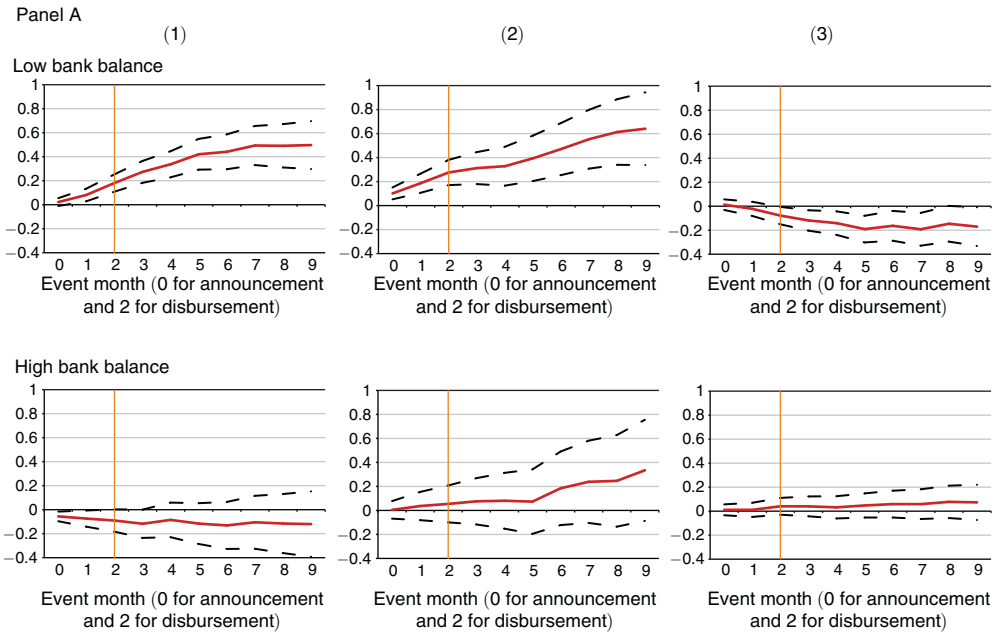


FIGURE 3. HETEROGENEITY IN SPENDING AND DEBT RESPONSE ACROSS CONSUMERS

(Continued)

across different financial instruments. In the following subsections, we estimate equation (4) for each group of comparison of consumers. To illustrate, if we have two groups ($N = 2$), we estimate each group's response to the stimulus in one regression using the specification in equation (4) (with group $N = 2$ being absorbed).

Since we have shown no pretreatment effect in Tables 2–3 and Figure 2, we absorb the pretreatment variable ($\$D_i \times 1_{pre}$) in the heterogeneity analysis to increase power and facilitate interpretation.¹⁹ We plot, for each group of consumers, the cumulative response coefficients (starting from the announcement month, $s = 0$), b_s , $s = 0-9$, along with their corresponding 95 percent confidence intervals (Figure 3). The remainder of the section focuses on the heterogeneous response between constrained and less constrained consumers based on two different proxies of constraints. We will discuss more heterogeneity in consumption and debt responses in the online Appendix.

Liquid Assets: Low Checking Account Balance versus High Checking Account Balance.—We first classify consumers based on the level of their liquid assets. A consumer is considered to have low liquid assets in our sample if his average monthly checking account balance in the four months before our analysis sample (i.e., 2010:04–2010:07) is below the twenty-fifth percentile of the distribution, or equivalently SG\$1,840 in the cross section of consumers in that period. Consumers

¹⁹We have conducted and verified that our results are qualitatively the same by adding the pretreatment months in the analysis. Specifically, the treatment and control groups exhibit parallel trends in the pretreatment period in all subgroups in the heterogeneity analysis.

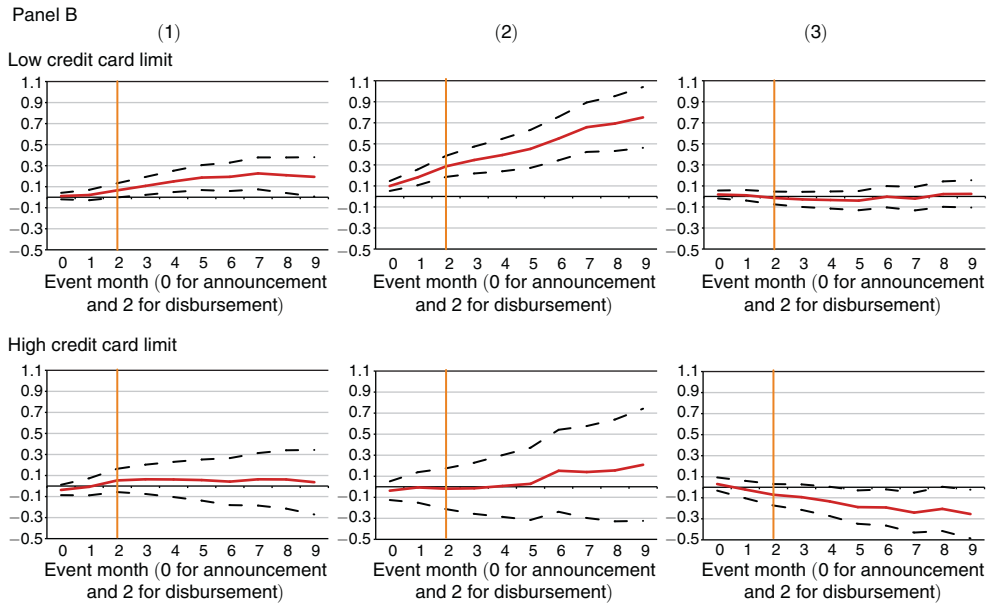


FIGURE 3. HETEROGENEITY IN SPENDING AND DEBT RESPONSE ACROSS CONSUMERS (Continued)

Notes: This figure plots the entire paths of cumulative coefficients $b_{s, s} = 0-9$, along with their corresponding 95 percent confidence intervals, across different consumers. The sample includes the matched treatment and control groups during the period of 2010:08–2011:11. For each comparison panel, column 1 shows the cumulative debit card spending response, column 2 shows the cumulative credit card spending response, and column 3 shows the cumulative credit card debt change response. Panel A compares consumers with low bank checking account balances (i.e., average checking account balance between 2010:04 and 2010:07 \leq SG\$ 1,840, or twenty-fifth of sample) with consumers with high bank checking account balances (i.e., average checking account balance between 2010:04 and 2010:07 \geq SG\$22,346, or seventy-fifth of the sample). Panel B compares consumers with high credit card limit (i.e., max credit card limit between 2010:04 and 2010:07 \geq SG\$ 9,000, or seventy-fifth of the sample) with consumers with low credit card limits (i.e., max credit card limit between 2010:04 and 2010:07 \leq SG\$5,000, or twenty-fifth of sample). The x -axis denotes the s th month after the announcement of the Growth Dividend Program, and the y -axis shows the dollar response (for every dollar received).

have high liquid assets if their average monthly balance in that period is above the seventy-fifth percentile of the distribution, or SG\$22,346. Consumers with low checking account balances are likely to be more liquidity constrained, as they cannot run down on their liquid assets in response to a negative shock.²⁰

Panel A of Figure 3 shows the comparison in the response path between these two groups of consumers. Low-balance consumers responded strongly via debit card spending: for each dollar of stimulus payment, $b_9 = \$0.50$ for low-balance consumers, and the effect is statistically significant (p -value < 0.001). In particular,

²⁰ On the other hand, high-balance consumers are also likely associated with precautionary saving motives. According to buffer stock models, consumers who face income uncertainty and incomplete insurance markets choose to delay consumption and save to prepare for negative income shocks in the future, especially in the presence of an incomplete market. Upon an unanticipated income shock, these consumers should respond by increasing their spending because the (unanticipated) positive income shock alleviates their income uncertainty. High balance is suggestive of the presence of a precautionary saving motive, but the measure may also reflect other offsetting factors (such as large wealth). Therefore, we study, within the high balance subsample, consumers who have low spending history during the out-of-sample history (i.e., 2010:04–2010:07). They are more likely to have precautionary saving motives given their high savings but low usage of their credit card limit. Results in the online Appendix (Figure A2) show a strong consumption response among these consumers, especially during the announcement period. This provides additional evidence to explain the announcement effect we document in Section IVB.

low-balance consumers started spending on debit cards from the second month of the announcement period and continued to experience a significant increase until month seven, after which the debit card spending increase plateaued. There is also a strong cumulative increase in credit card spending among low-balance consumers: $b_9 = \$0.64$ for each dollar received, which is statistically significant (p -value < 0.001). On the other hand, high-balance consumers did not increase debit card spending, and their cumulative credit card spending increase by month nine is equal to $\$0.33$ per dollar received, which is smaller than that of the low-balance consumers and is statistically insignificant. We perform a formal test of the difference in total cumulative spending response between low- and high-balance consumers. We run an F -test of the difference in the cumulative coefficients b_9 between the two groups. The result is statistically significant (p -value < 0.001), indicating that low-balance consumers spent more via debit cards and credit cards than high-balance consumers.

Low-balance consumers started to pay down their credit card debt upon receipt of the stimulus money (month two after the announcement), and by month nine, the cumulative debt decrease is $\$0.17$ per dollar received and is statistically significant (p -value = 0.038). In contrast, there is no change in credit card debt among high-balance consumers.

Credit Access: High Credit Card Limit versus Low Credit Card Limit.—We then study the heterogeneous response by consumers with different credit access. We classify consumers in our sample as having a high credit card limit if their maximum credit card limit in the four months before our analysis sample (i.e., 2010:04–2010:07) is above the seventy-fifth percentile of the distribution, or equivalently SG\$9,000 in the cross section of consumers during that period. Consumers have a low credit card limit if their maximum credit card limit between 2010:04 and 2010:07 is below the twenty-fifth percentile of the sample, or SG\$5,000. This is another measure to capture credit constrained consumers that has been used in previous studies. Consumers with low credit card limits presumably have limited access to the credit market, making it difficult for them to borrow (using credit cards) to smooth consumption.

Panel B of Figure 3 shows the comparison across these two groups of consumers. High credit card limit consumers showed little spending response, regardless of the type of financial instruments. The cumulative spending coefficients for both credit card and debit card are statistically insignificant throughout the period. Low credit card limit consumers reacted to the stimulus program by increasing both their debit card and credit card spending. However, the effect is stronger on credit card spending. The cumulative debit card spending increase at month nine after the announcement month is $b_9 = \$0.19$ per dollar received (p -value = 0.042). Credit card spending has a cumulative increase of $\$0.75$ per dollar received by month nine (p -value < 0.001). An F -test of the cumulative coefficients of total spending suggests that low credit limit consumers' total spending response is greater than that of high credit limit consumers (difference = $\$0.70$ with p -value = 0.02).

While low credit limit card consumers saw no credit card debt change during the ten-month period, high credit card limit consumers' credit card debt decreased strongly: by month nine, the cumulative credit card debt change is $-\$0.25$ per $\$1$ received (p -value = 0.031).

Using both bank account balance and credit card limit, we document results that are consistent with the literature (Gross and Souleles 2002; Agarwal, Liu, and Souleles 2007): constrained consumers react strongly to the stimulus in spending, and they also use the positive income shock to reduce their credit card debt. In addition, results in both subsample analyses also reveal the mechanism through which constrained consumers (rationally) responded to the income shock. Specifically, the dynamics in consumption response we uncover in Section IVD (that consumers used credit card first, followed by debit card after disbursement, and finally credit card) are largely driven by constrained consumers, i.e., those with low bank balance or low credit card limit. During the announcement period, consumers used the credit market to respond to the unanticipated income shock and smooth their consumption. Although theory defines liquidity constrained consumers as those who strictly cannot borrow or have zero liquid wealth, we identify constrained consumers in our dataset based on the relative level of savings or credit access. As a result, our identified constrained consumers have some access to credit, and will use this costly means to smooth consumption upon announcement. After disbursement, (constrained) consumers started to use their debit card more, consistent with the fact that debit card spending is a cheaper way to consume as the realized income increase relaxes their constraints. Debit card is cheaper than credit card because of the possibility that consumers can run into (costly) credit card debt by using credit cards, given the fact that they are *ex ante* constrained. Subsequently, as the stimulus money was being used up, constraints start to become binding again, leading them to switch to credit cards to increase spending. Overall, this spending pattern among the constrained consumers suggests that financial incentives drive their behavior.

B. Heterogeneity in Spending Response: by Spending Category

The extant literature documents heterogeneity in the type of spending response to positive income shocks (e.g., Parker et al. 2013). In our data, merchant type descriptions are provided in the debit and credit card transactions, from which we group them into the following eight categories: supermarket, service, dining, entertainment, apparel, travel, small durable goods, and online. We leave the detailed description of the analysis in the online Appendix. To summarize, we find that discretionary spending categories, such as apparel and travel, responded strongly to the stimulus program. Consistent with the existing literature (Parker et al. 2013), consumption also responded significantly in the small durable goods category. Although consumers are unlikely to increase car or house consumption after the Growth Dividend Program, they appeared to increase spending on less costly durable goods such as electronics, computers, home or office furnishings, and appliances. Consumers also increased their spending in other categories such as supermarket, dining, entertainment, and transportation, but economically and statistically the effect is weaker.

C. Robustness

We perform additional tests to study the robustness of our results. First, to further address the concern that foreigners differ from Singaporeans in unobservable ways that may affect their spending behavior, we (i) restrict the control group to

foreigners of certain nationalities that either come from neighboring countries or have ethnic and cultural backgrounds similar to Singaporeans, and (ii) completely drop the foreigners from the sample and perform tests by exploiting the heterogeneity in the payout amount. Specifically, we use the Singaporeans with the smallest amount of the growth dividend (i.e., those with an annual income greater than SG\$100,000) as the control group. Second, we restrict our treatment group to those unaffected by other potential treatments at the same time (e.g., other government programs, and the annual bonus payout). Last, we investigate the robustness of our statistical inference—consistency of standard errors—and conduct our tests using alternative specifications as suggested by the literature (Bertrand, Duflo, and Mullainathan 2004; Abadie and Imbens 2006). Throughout the robustness tests, our main findings remain the same and we leave the details to the online Appendix.

V. Conclusion

This paper uses a unique, new panel dataset of credit card, debit card, and checking account information for more than 180,000 consumers in Singapore to analyze how consumption and debt responded to an unanticipated fiscal stimulus program announced on February 18, 2011. The unique policy experiment by the Singapore government allows us to distinguish (or to estimate) an announcement effect and a disbursement effect. We use a difference-in-differences identification to estimate the month-by-month response to the program. Foreigners were not eligible for the growth dividend; this exclusion restriction allows us to identify the causal effect of the program on spending by using foreigners as our control group.

We find that consumption rose significantly after the fiscal policy announcement: for each dollar received, consumers on average spent \$0.80 (aggregated across different financial accounts) during the ten months after the announcement. Consumers' credit card debt moderately decreased during this period. We identify a strong announcement effect: consumers started to increase spending during the two-month announcement period before the cash payout. We also find that consumption response is distributed across debit card spending (25 percent of the total response) and credit card spending (75 percent of the total response). More importantly, consumers started spending via credit cards during the announcement period, then switched to debit cards after disbursement, before finally significantly increasing their credit card usage. There is significant heterogeneity in the response to the fiscal stimulus. Consumption rose primarily in the discretionary or small durable goods categories. Consumers with low liquid assets or with a low credit card limit showed a strong consumption response. In comparison with previous findings (e.g., Agarwal, Liu, and Souleles 2007), our MPC estimates are larger for two main reasons. We are able to capture the announcement effect and the spending on debit cards during the disbursement period. Excluding these two effects, our estimates are consistent with Agarwal, Liu, and Souleles (2007). This difference has significant implications from a public policy perspective: if the goal of the policy is to maximize consumption response, then the past studies underestimate the effect due to lack of data (debit cards) and design of policy (no identification for the announcement effect).

Our main contributions in relation to prior literature are fourfold. First, we are the first to document (a significant) announcement effect of this stimulus program,

which is consistent with life-cycle theory predictions on consumption response to unanticipated income shocks. Second, the decomposition of the consumption response into different spending instruments implies that prior work based on micro-data from one single payment instrument (e.g., credit card) likely underestimates the spending response to income shocks. Third, we document the dynamics of consumption response across different spending instruments—a rise in credit card spending following the program’s announcement, then the switch to (cheaper) debit card spending after the disbursement of the stimulus, and finally the switch back to credit card spending in the later months. This newly documented consumption response mechanism highlights that financial incentives drive (constrained) consumers’ spending behavior. Finally, our data richness allows us to identify various measures of (financial) constraints, and our findings provide additional credibility to the result that constraints—whether labeled credit or liquidity constraints—are important for consumption response to income shocks.

REFERENCES

- Aaronson, Daniel, Sumit Agarwal, and Erik French. 2012. “Spending and Debt Response to Minimum Wage Hikes.” *American Economic Review* 102 (7): 3111–39.
- Abadie, Alberto, and Guido W. Imbens. 2006. “Large Sample Properties of Matching Estimators for Average Treatment Effects.” *Econometrica* 74 (1): 235–67.
- Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles. 2007. “The Reaction of Consumer Spending and Debt to Tax Rebates—Evidence from Consumer Credit Data.” *Journal of Political Economy* 115 (6): 986–1019.
- Agarwal, Sumit, and Wenlan Qian. 2014. “Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore: Dataset.” *American Economic Review*. <http://dx.doi.org/10.1257/aer.104.12.4205>.
- Baker, Scott R. 2014. “Debt and the Consumption Response to Household Income Shocks.” Unpublished.
- Bertrand, Marianne, Esther Dufo, and Sendhil Mullainathan. 2004. “How Much Should We Trust Differences-in-Differences Estimates?” *Quarterly Journal of Economics* 119 (1): 249–75.
- Browning, Martin, and Thomas F. Crossley. 2001a. “The Life-Cycle Model of Consumption and Saving.” *Journal of Economic Perspectives* 15 (3): 3–22.
- Browning, Martin, and Thomas F. Crossley. 2001b. “Unemployment Insurance Benefit Levels and Consumption Changes.” *Journal of Public Economics* 80 (1): 1–23.
- Carroll, Christopher D. 1992. “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence.” *Brookings Papers on Economic Activity* 23 (2): 61–156.
- Carroll, Christopher D. 1997. “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis.” *Quarter Journal of Economics* 112 (1): 1–55.
- Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2014. “Harnessing Naturally Occurring Data to Measure the Response of Spending Income.” *Science* 345 (6193): 212–15.
- Gertler, Paul, and Jonathan Gruber. 2002. “Insuring Consumption Against Illness.” *American Economic Review* 92 (1): 51–76.
- Gross, David B., and Nicholas S. Souleles. 2002. “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data.” *Quarterly Journal of Economics* 117 (1): 149–85.
- Gruber, Jonathan. 1997. “The Consumption Smoothing Benefits of Unemployment Insurance.” *American Economic Review* 87 (1): 192–205.
- Ho, Teck-Hua, Ivan P. L. Png, and Sadat Reza. 2014. “Sunk Cost Fallacy in Driving the World’s Costliest Cars.” Unpublished.
- Hsieh, Chang-Tai. 2003. “Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund.” *American Economic Review* 93 (1): 397–405.
- Jappelli, Tullio, and Mario Padula. 2014. “The Consumption and Wealth Effects of an Unanticipated Change in Lifetime Resources.” University of Naples, Centre for Studies in Economics and Finance (CSEF) Working Paper 354.

- Jappelli, Tullio, Jörn-Steffen Pischke, and Nicholas S. Souleles.** 1998. "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources." *Review of Economics and Statistics* 80 (2): 251–62.
- Jappelli, Tullio, and Luigi Pistaferri.** 2010. "The Consumption Response to Income Changes." *Annual Review of Economics* 2: 479–506.
- Jappelli, Tullio, and Luigi Pistaferri.** 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics* 6(4): 107–36.
- Johnson, David S., Jonathan A. Parker, and Nicholas Souleles.** 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96 (5): 1589–1610.
- Leth-Petersen, Søren.** 2010. "Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit?" *American Economic Review* 100 (3): 1080–1103.
- Parker, Jonathan A.** 1999. "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes." *American Economic Review* 89 (4): 959–73.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland.** 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6): 2530–53.
- Paxson, Christina H.** 1993. "Consumption and Income Seasonality in Thailand." *Journal of Political Economy* 101 (1): 39–72.
- Poterba, James M.** 1988. "Are Consumers Forward Looking? Evidence from Fiscal Experiments." *American Economic Review* 78 (2): 413–18.
- Shapiro, Matthew D., and Joel Slemrod.** 1995. "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding." *American Economic Review* 85 (1): 274–83.
- Shapiro, Matthew D., and Joel Slemrod.** 2003a. "Consumer Response to Tax Rebates." *American Economic Review* 93 (1): 381–96.
- Shapiro, Matthew D., and Joel Slemrod.** 2003b. "Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys." In *Tax Policy and the Economy*, edited by James M. Poterba, 83–109. Cambridge, MA: MIT Press.
- Souleles, Nicholas S.** 1999. "The Response of Household Consumption to Income Tax Refunds." *American Economic Review* 89 (4): 947–58.
- Souleles, Nicholas S.** 2000. "College Tuition and Household Savings and Consumption." *Journal of Public Economics* 77 (2): 185–207.
- Souleles, Nicholas S.** 2002. "Consumer Response to the Reagan Tax Cuts." *Journal of Public Economics* 85 (1): 99–120.
- Stephens, Melvin Jr.** 2003. "'3rd of the Month': Do Social Security Recipients Smooth Consumption between Checks?" *American Economic Review* 93 (1): 406–22.
- Stephens, Melvin Jr.** 2006. "Paycheque Receipt and the Timing of Consumption." *Economic Journal* 116 (513): 680–701.
- Stephens, Melvin Jr.** 2008. "The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans." *Review of Economics and Statistics* 90 (2): 241–52.
- Stephens, Melvin Jr., and Takashi Unayama.** 2011. "The Consumption Response to Seasonal Income: Evidence from Japanese Public Pension Benefits." *American Economic Journal: Applied Economics* 3 (4): 86–118.
- Wolpin, Kenneth.** 1982. "A New Test of the Permanent Income Hypothesis: The Impact of Weather on the Income and Consumption of Farm Households in India." *International Economic Review* 23 (3): 583–94.
- Zeldes, Stephen P.** 1989a. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97 (2): 305–46.
- Zeldes, Stephen P.** 1989b. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics* 104 (2): 275–98.