HARVARD | BUSINESS | SCHOOL

Food Stamp Entrepreneurs

Gareth Olds

Working Paper 16-143

Food Stamp Entrepreneurs

Gareth Olds
Harvard Business School

Working Paper 16-143

FOOD STAMP ENTREPRENEURS

Gareth Olds*

This draft: May 2016

Abstract

This paper explores how eligibility for the Supplemental Nutrition Assistance Program (SNAP, formerly known as the Food Stamps Program) affects firm formation. Using a variety of identification strategies, I show that expanded SNAP eligibility in the mid-2000s increased enrollment by 3-5 percentage points. Newly-eligible households were also 20% more likely to own a business, with larger effects for incorporated firms. I find large increases in labor supply on the extensive and intensive margins, equivalent to 1.1 million additional workers. I also develop a series of falsification checks that use information from unaffected portions of the income distribution to improve RD estimates. This strategy improves balance on observables between treatment and control groups, and Monte Carlo simulations find a significant reduction in Type 1 Error. Finally, I show that the empirical results are driven entirely by newly-eligible households that did not enroll, suggesting uninsured risk from leaving wage employment is the primary barrier to entrepreneurship for this population.

*Harvard Business School. Email: golds@hbs.edu. This research was supported in part by the National Science Foundation's IGERT Fellowship and the Hazeltine Fellowship for Research in Entrepreneurship. I am deeply grateful to Ken Chay, Brian Knight and David Weil for their feedback and support. A special thanks to Alex Eble, Alexei Abrahams, Andrew Elzinga, Bruno Gasperini, Shiva Koohi and Tim Squires for useful suggestions.

1 Introduction

What is the central barrier to firm formation? Do most households face credit rationing, or is uninsured risk from leaving wage employment the binding constraint? Early attempts to answer this question have focused on unexpected income shocks, such as inheritances. However, since only 16% of adults plan to leave a substantial inheritance to their children, and the unpredictability of these windfalls makes business planning difficult, it is not clear how relevant inherited wealth is for the representative aspiring entrepreneur.[1]

Public programs, on the other hand, offer a predictable income stream that provide support to millions of Americans. The Supplemental Nutrition Assistance Program (SNAP, formerly known as the Food Stamps Program) currently has 50 million enrolled members, providing income to nearly one in six Americans and a fall-back option for many more. If credit and risk constraints are important restrictions for potential entrepreneurs, a natural question is whether generous public insurance systems allow people to start firms.

This paper identifies the effect of SNAP eligibility on business ownership. I examine an expansion of SNAP threshold levels beginning in the mid-2000's to determine the impact on rates of self-employment and birth of incorporated firms. I use three identification strategies to isolate this relationship: first, I set up a simple difference-in-differences framework using cross-state variation; second, I use a regression discontinuity framework that incorporates pre-policy data as a falsification check (see Olds, 2014 for a formal description); and finally, I present a falsification exercise that partitions the distribution of the running variable and controls for misspecification bias arising from non-linearity in the conditional expectation. This strategy is applicable to a general class of RD designs, requires no additional data, and outperforms the traditional RD estimator in Monte Carlo simulations for a range of performance metrics and bandwidths values.

Using these strategies I find that becoming eligible for SNAP increased a household's likelihood of enrollment by nearly 3-5 percentage points. Newly-eligible households were also around 20% more likely to own a business. Most of the increase in self-employment is the result of new firm birth—which grew by 12%—rather than increased firm survival. These tend to be high-quality firms: the marginal effects were particularly strong for incorporated ventures, with the probability of owning an incorporated business increasing by 16% as a result of the policy. The expansion of SNAP also increased the length of the work-year by 2.5% and the work-week by 5% relative to the baseline, a labor supply increase equivalent to 19.6 million full-time workers. There is little evidence that observable characteristics differ between treatment and control groups using the partitioned RD

[1] According to a T. D. Ameritrade study: http://amtd.com/newsroom/releasedetail.cfm?releaseid=702709.

Table 1: Motivating Facts

Business Ownership	Frequency	Households that are ... Below 200% FPL	Receiving Medicaid	Receiving SNAP	On public assistance
None	88.3%	32.8%	17.8%	8.1%	24.2%
Unincorporated	7.8%	26.4%	14.8%	4.4%	19.6%
Incorporated	4.0%	8.3%	7.3%	0.8%	10.6%

Data: CPS, 1996-2011. "FPL" is the Federal Poverty Line, using HHS guidelines by year/household size. "Public assistance" includes SNAP, Medicaid, SCHIP and unemployment benefits.

method. Finally, the results are driven entirely by newly eligible non-enrollees. Taken together, these results strongly suggest the presence of a large population of would-be entrepreneurs held back by uninsured risk.

The following section provides some background motivation and literature on the link between public assistance and entrepreneurship, discusses the recent expansion of SNAP, and describes the data. Section 3 presents the difference-in-difference and regression discontinuity results. Section 4 explores firm origins and characteristics, as well as some potential mechanisms. Section 5 formally describes the partitioned procedure and presents the results. Section 6 discusses several criteria for model validation: balance of observable characteristics between treatment and control groups, falsification and robustness tests, and Monte Carlo simulations evaluating the behavior of various models. Section 7 concludes.

2 Background

Do entrepreneurs use public programs? Table 2 breaks down US households by whether they report having an unincorporated business, an incorporated one, or neither. Twelve percent of households have a working adult classified as "self-employed", and 30% of those—4% of the total population—have an incorporated business. To give a sense of each group's income distribution, Table 2 also reports the proportion of each group below 200% of the Federal Poverty Line, which varies by household size. Around one-third of non-business-owning households have income below this level, compared with just over one-quarter for households with unincorporated businesses. Even households with incorporated firms, who are typically thought of as more "serious" entrepreneurs (Levine and Rubinstein, 2012), have a one-in-twelve chance of being below this level.

Also shown are the rates of public program usage by these groups. First, note that usage by households with firms is strictly non-zero: more than one in ten households with an incorporated firm use SNAP, Medicaid, SCHIP (a child health insurance program) or unemployment benefits. Second, uptake of some programs is disproportionately high for entrepreneurs considering their income distribution. Non-business-owning households are 4 times more likely to be below 200% of the poverty line than households with incorporated firms, yet they are less than three times more likely to use Medicaid and only about twice as likely to be on any form of public assistance. Entrepreneurs tend to have slightly higher income than households with only wage income, yet take advantage of public programs at higher rates than would be predicted by their income.

Literature.

There is a sizable literature devoted to credit rationing and whether it restricts the growth of entrepreneurial ventures. There are two types of credit rationing relevant to new businesses: Type I rationing—when a loan applicant receives a smaller loan he or she would like—and Type II rationing—failure of some people to get a loan altogether. Theoretical treatments of credit rationing emphasize the importance of bankruptcy and monitoring costs, asymmetric information about heterogeneous borrowers, and uncertainty about returns (Parker, 2009). One testable implication of rationing is that entry into entrepreneurship is positively related to existing assets.[2]

Early tests of this relationship provide evidence that credit rationing occurs. Evans and Jovanovic (1989) find a positive and significant effect of assets on entrepreneurship. They argue that 94% of people who are likely to start a business have faced credit rationing, and that 1.3% of the US population have been prevented from starting a business as a result. Concerned about the potential endogeneity of wealth, later research focused on windfalls such as inheritances. Holtz-Eakin et al. (1994a), Holtz-Eakin et al. (1994b), and Blanchflower and Oswald (1998) all find a strong relationship between inherited funds and entrepreneurship, consistent with credit market imperfections that prevent potential entrepreneurs from borrowing against a firm's future income stream.

There are also a number of papers in development economics on the importance of credit constraints to business success (see Besley, 1995 for an overview and Karlan et al., 2012 for a recent example). However, most of these focus on farmers, whose constraints may not be representative of other enterprises. The study most closely related to the

[2]The monotonic relationship between assets and entrepreneurship has been called into question by more recent research by Buera (2009).

topic of this paper is Bianchi and Bobba (2013), which provides evidence that *Progresa* increased entrepreneurship in Mexico. If entrepreneurs in the developing world face a different set of constraints than in advanced economies, however, than it is difficult to generalize these results to the industrialized world.

The credit rationing literature has some limitations. Most of the research designs used in the literature are not quasi-experimental or connected to a policy, raising concerns about selection on unobservables. In particular, it could be the case that unobservable, inherited characteristics make certain people more likely to be entrepreneurs, and these same characteristics are associated with wealth accumulation and bequests. Individuals entering self-employment after receiving an inheritance could be pursuing an occupational choice strategy in which bequests are simply a cheap source of start-up funds relative to costly debt financing. These individuals need not be credit constrained: they may have been perfectly capable of receiving a loan, but saw no need to apply when they anticipated an inheritance in the near future. Ameliorating credit rationing would not facilitate entrepreneurship in these cases.

This paper uses a different strategy, focusing on exogenous access to income support provided by changes in SNAP eligibility.[3] If credit rationing stops potential entrepreneurs from borrowing to finance their businesses, receipt of public benefits should allow more people to enter self-employment. Even for households that do not receive benefits, however, the presence of an income floor that comes with SNAP eligibility reduces the riskiness of leaving wage employment. Since self-employment is an inherently risky endeavor, these households may be more likely to leave a wage-earning job when given access to insurance against consumption shocks. Disentangling these two effects—risk versus credit—is difficult, but I show some preliminary evidence later in the paper that suggests that uninsured risk is the binding constraint in firm formation, which is consistent with other studies (Karlan et al., 2012; Bianchi and Bobba, 2013).

Policy history.

Welfare reform during the 1990s, exemplified by the Personal Responsibility and Work Opportunity Reconciliation Act of 1996, sharply reduced the number of people who received government assistance. While reform focused primarily on Aid to Families with

[3]Technically SNAP is not an income support program because it provides funds that can only be spent on food, as opposed to Temporary Assistance for Needy Families (TANF, often called “welfare”) funds, which can be used for general purchases. However, recent evidence suggests most recipient households are infra-marginal in food consumption, so that the impact of SNAP benefits on household spending is the same as a similarly-sized cash transfer (Hoynes and Schanzenbach, 2009).

Dependent Children (AFDC, often referred to as "welfare" and later called Temporary Assistance for Needy Families), other social programs such as Food Stamps saw new time limits, work requirements and deduction caps, pushing down SNAP enrollment to a low of 17 million by the end of the decade.

However, a November 2000 clarification of the SNAP federal guidelines gave states more flexibility in implementation. The new interpretation of the statute allowed states to waive asset tests and raise the gross income limit to 200% of the Federal Poverty Line (FPL) when determining eligibility for SNAP, up from the previous level of 130 percent. This expansion was based on a loophole in which individuals could be considered "categorically eligible" if they met less restrictive requirements, so long as they previously received benefits from the state agency administering the program. Under the new interpretation, "benefits" included in-kind transfers, which in most states is satisfied by being handed a free pamphlet or visiting a government website, since these services are paid for with public funds. States expanded their programs in a staggered fashion over the last decade, choosing a variety of thresholds between 130% and 200% of the FPL.[4] Since the rule change in 2000, SNAP participation rates have more than doubled, and enrollment climbed to an all-time high of nearly 47 million people in 2012.

Data.

To understand whether this expansion of the Food Stamp program encouraged entrepreneurship by relaxing credit constraints, I use household-level data from the Current Population Survey (CPS) March Supplement files for the years 1996-2011. Eligibility for SNAP is based on total household income at the time of application. Since states are required by law to provide benefits to eligible households no later than 30 days from the application date,[5] I use the income reported by a household in a given year to assess eligibility. Summary statistics can be found in the Appendix.

Treatment status in this context means new eligibility under the change in rules. Households will be "treated" if their income is below the new eligibility threshold—the threshold percentage of the poverty line times the FPL, adjusted for household size and year—and greater than 130% of the FPL. Households below 130% FPL are not included in this study, which focuses the analysis only on people who are either newly eligible or ineligible rather than those who have always been eligible. State threshold levels and dates of policy implementation come from Mathematica Policy Research's comprehensive report on SNAP

[4] However, once a policy threshold was chosen, nearly all states kept the cut-off in place without revision.
[5] See 7 USC 2020(e)(3); in come cases benefits must begin as soon as 7 days from the application date.

eligibility laws.[6]

3 Basic Results

Difference-in-Differences.

Self-employment rates and SNAP enrollment are both likely driven by underlying household characteristics, such as wealth, education, and access to credit. In particular, households are eligible only if their income is low enough; entrepreneurship often depends positively on income, so treated households will necessarily have lower rates of self-employment. A simple way to overcome these spurious drivers is to focus on differences in uptake and entrepreneurship between the treated and untreated groups, before and after the policy was enacted. So long as the outcomes for treated and untreated households would have changed in relatively the same way in the absence of the policy—the parallel trends assumption—the difference-in-differences estimator will identify the causal effect of the program expansion.

Figure 3 illustrates this logic by plotting average SNAP enrollment and self-employment against the number years until the policy took effect in a given state. Enrollment in SNAP is non-zero both for the ineligible group and the newly-eligible population before the rule change, suggesting measurement error in household income, flexibility in enrollment beyond the federal rules, or a combination of the two. However, there is a sharp increase in uptake for the newly-eligible group after the change, whereas ineligible household enrollment remains flat. Similarly, self-employment rates are lower for the newly eligible population but relatively flat before the policy shift and then begin to rise after the change.

Translating these results into a regression framework is straightforward. In this set-up,

$$\begin{gathered} Y_{ist} = \beta_0 + \beta_1 \cdot Treat_{it} Post_{st} + \beta_2 \cdot Treat_{it} + \beta_3 \cdot Post_{st} \\ \xi \cdot X_{it} + \nu_s + \eta_t + \gamma \cdot t\nu_s + \varepsilon_{ist} \end{gathered} \tag{1}$$

for individual i in state s at time t, where Y_{ist} is either SNAP enrollment or self-employment. Treatment status $Treat_{it} = \mathbf{1}[Inc_{it} \leq Thresh_{st}]$ and uses a household's combined income to determine eligibility; this variable is only defined for households above the previous cut-off of 130% FPL. The variable $Post_{st} = \mathbf{1}[t \geq PolicyYear_s]$, indicating whether an observation is before or after the policy's enacting. The parameters ν_s and η_t are state

[6] "Non–Cash Categorical Eligibility For SNAP: State Policies and the Number and Characteristics of SNAP Households Categorically Eligible Through Those Policies", available at mathematica-mpr.com/publications/PDFs/nutrition/non-cash_snap.pdf.

Figure 1: Difference-in-Differences

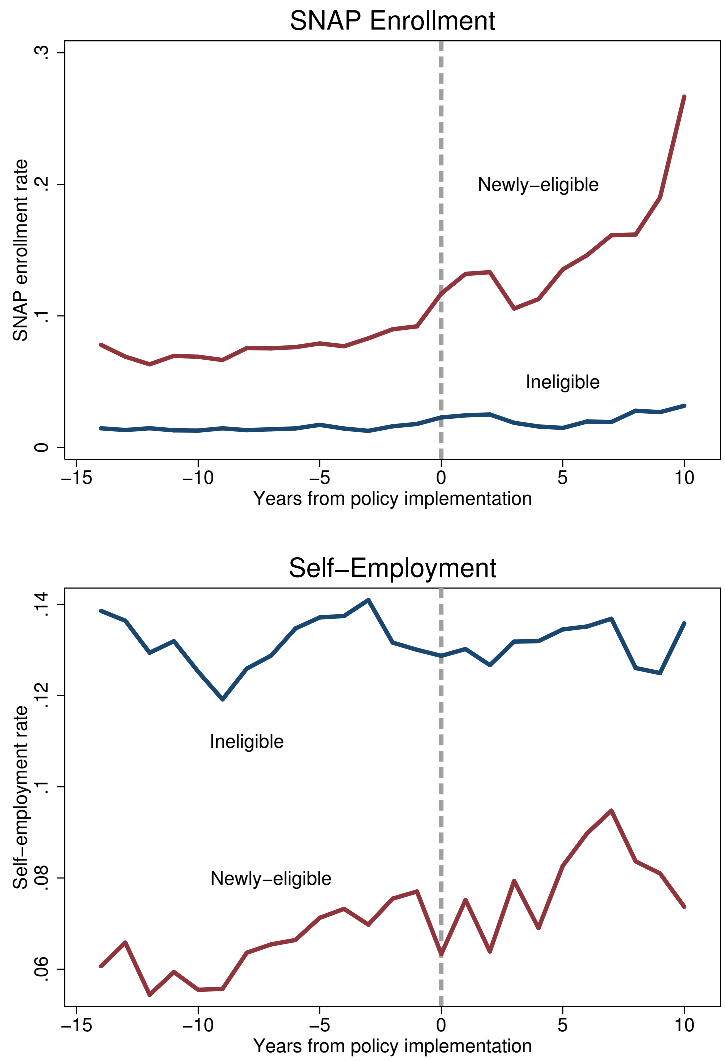

Table 2: Difference-in-Differences

	SNAP			Self-Employment		
	(1)	(2)	(3)	(4)	(5)	(6)
$Treat \cdot Post$	0.0510**	0.0481**	0.0389**	0.0104*	0.0128**	0.0101+
	(0.009)	(0.008)	(0.008)	(0.004)	(0.004)	(0.006)
$Treat$	0.0655**	0.0490**	0.0302**	-0.0653**	-0.0671**	-0.0175**
	(0.005)	(0.004)	(0.003)	(0.004)	(0.003)	(0.004)
$Post$	-0.0041*	-0.0033+	-0.0026	-0.00167	-0.0036+	-0.0028
	(0.002)	(0.002)	(0.002)	(0.004)	(0.002)	(0.003)
$SNAP_{t-1}$			0.2790**			-0.0043+
			(0.01)			(0.003)
Pre-policy avg.	0.021	0.021	0.021	0.126	0.126	0.126
ATE	245.9%	232.1%	187.7%	8.3%	10.2%	8.1%
Treated avg.	0.077	0.077	0.077	0.067	0.067	0.067
TOT	65.9%	62.2%	50.3%	15.6%	19.2%	15.2%
Covariates	No	Yes	Yes	No	Yes	Yes
Observations	603,058	600,118	164,685	603,058	603,058	164,685
R-squared	0.032	0.074	0.166	0.003	0.008	0.044

OLS. All specifications include time/state FE and state trends. Data: CPS, 1996-2011.
** 0.01, * 0.05, + 0.1. Robust standard errors clustered at the state level in parentheses.

and year fixed effects, $\gamma \cdot t\nu_s$ is a state-specific linear time trend, and X_{it} is a vector of covariates. The difference-in-difference estimator of the effect of the expansion is β_1.

Table 3 shows the results of this specification, including the average treatment effect relative to the pre-policy baseline for the entire population, and the treatment on the treated, which uses the pre-policy average of the treated group as the denominator. Columns (1)-(3) use SNAP enrollment as the dependent variable; since the data are aggregated to the household level, this variable takes a value of one if any members of the household are receiving Food Stamps. Columns (4)-(6) use self-employment status, which is one whenever a member of the household reports self-employment income. The results show a 5 percentage point increase in receipt of SNAP benefits, with the coefficient dropping to 4.8 percentage points covariates are included (see the appendix for a complete list of variables). Marginal effects are somewhat misleading, since none of these households were previously eligible even though around 7% were enrolled in SNAP before the policy change. However, the level is consistent with the 6 percentage point increase in enrollment between 2000 and 2011.

Concurrent with this increase in Food Stamp use was a 1 percentage point rise in self-employment for treated individuals in the post period, relative to untreated and pre-policy households. This result means that expanded SNAP eligibility increased the number of households engaged in entrepreneurship by around 8-10% in the total population and 15-20% among newly eligible households.

Problems.

Difference-in-difference strategies have some well-known issues, particularly in the presence of serial correlation (see Bertrand et al., 2004). If the changes identified by the difference-in-difference estimator are solely the result of the policy, underlying observable and unobservable characteristics between the two groups should not change concurrently. While differential changes in unobservables are assumed away as part of the parallel trends assumption, the shift in observables can be observed directly. Column 1 of Table 6 repeats the specification from Table 3 but each of the covariates as the outcome variables. Ten of these 18 covariates are significantly different at the 5% level for the treated group after the policy took effect, and it is unlikely these changes are the result of the expansion in SNAP eligibility, particularly for demographic variables. If changes in the treated and untreated group were random, only 5% of these variables would show up as significant, rather than the 56% observed. The fact that observable characteristics differ so starkly makes it unlikely that unobservables are balanced between these groups, casting doubt on the argument that β_1 isolates the causal effect of the expansion.

Regression Discontinuity with a Pre-Policy Falsification Check

One way to address this imbalance in observable characteristics is to take advantage of information about eligibility thresholds to exploit discontinuities in enrollment rates. Eligibility for SNAP is determined by income: people just below the threshold can receive SNAP benefits, whereas those below cannot. People on either side of the threshold are more likely to be similar than those far away, with the crucial difference that those below are marginally more likely to receive Food Stamps. If people are unaware of the policy cut-off levels before the policy was expanded—or are otherwise unable to sort into or out of treatment by manipulating their income—SNAP eligibility is close to random for households near the threshold. A frequently-used test for whether households are sorting into treatment is to check whether the density of the population varies at the policy cut-off; any "clumping" could indicate manipulation of the variable that determines treatment (see McCrary, 2008). Figure A2.1 in the appendix shows a non-parametric estimate of the density of the forcing variable both before and after the policy was enacted, and there does not appear to be any clumping below the threshold in either period.

Comparing people to the left and the right of the threshold, an identification strategy known as regression discontinuity, can easily accommodate a falsification check using pre-policy data to control for underlying changes in the distribution of the running variable. This falsification test essentially estimates an RD in both the post-policy and pre-policy periods, identifying the treatment effect as the differences in the threshold breaks (see Olds, 2014 for a formal description of this procedure).

To estimate the model, I focus on two of the most popular methods for RD designs: high-order polynomials and local linear regression. In the context of the expansion of the Food Stamps program, this means estimating

$$Y_{ist} = \beta_0 + \beta_1 \cdot Treat_{it} \cdot Post_{st} + \beta_2 \cdot Treat_{it} + \beta_3 \cdot Post_{st} + f(Rule_{it}) + \nu_s + \eta_t + \gamma \cdot t\nu_s + \varepsilon_{ist} \quad (2)$$

for the polynomial regression, where $f(\cdot)$ is a high-order function of the forcing variable $Rule_{it} = Inc_{it} - Thresh_{ist}$. For local linear regression, this becomes

$$Y_{ist} = \beta_0 + \beta_1 \cdot Treat_{it} \cdot Post_{st} + \beta_2 \cdot Treat_{it} + \beta_3 \cdot Post_{st} + \beta_4 \cdot Rule_{it} + \nu_s + \eta_t + \gamma \cdot t\nu_s + \varepsilon_{ist}$$
$$\text{if } abs(Rule_{it}) \leq W \quad (3)$$

where W is the bandwidth around the threshold. In both cases, β_1 identifies the treatment effect at the discontinuity, $\tau_{ATE|X=c}$. I use a 6th-order function for $f(\cdot)$ and present a variety of bandwidths W for robustness.

Figure 2 provides some intuition for what the augmented RD estimator is picking up. The top panel shows the non-parametric pre-post difference in SNAP enrollment as a function of income, generated by pooling households into equally-sized bins by income and taking the within-bin difference between post- and pre-policy observations. The conditional expectation is normalized so that the eligibility threshold is zero. A traditional difference-in-difference set-up takes the average of the graph to the left of the cut-off minus the average of the graph to the right. Regression discontinuity on the differenced data, on the other hand, only uses observations that are close to the threshold, isolating the jump in the graph near zero. The intuition for this strategy is that differences in enrollment near the threshold are more likely to be driven by the policy—since having an income that is a few dollars above or below the cut-off may be random—than averages over a wider set of incomes—since these groups may be different in other ways. This procedure improves on a basic RD in the post period by incorporating a “falsification check” directly into the estimator, since the treatment effect nets out any enrollment jumps nears the break that pre-dated the policy.

The lower panel repeats the process for self-employment, which jumps upward at the threshold by about one percentage point. However, there are other changes elsewhere in the income distribution that are probably unrelated to the policy, so any estimator that picks up these differences is likely to be biased. The regression discontinuity design is identified using only people for whom treatment status is pseudo-randomly assigned. Changes in self-employment that occur among the same population are more convincingly driven by policy than changes across a wider subset of the data.

Table 3 shows the results of the RD design for SNAP enrollment, along with the local average treatment effect and the treatment on the treated. Panel A estimates equation (2) using a 6th-order polynomial for $f(\cdot)$ and equation (3) for three bandwidths W, using SNAP enrollment as the dependent variable; panel B does the same for self-employment. Using the polynomial method, the eligibility expansion pushed up enrollment by 4.8 percentage points, which is a 62% increase from the treated group baseline of 7.7%. The bandwidth method reduces this number to a 2.3 percentage point increase, an enrollment growth rate of 30% among the treated.

Self-employment rates rose by 1.8 percentage points using the polynomial method, which is a 28% increase for the treated population (from the pre-policy average of 6.7%) or a 15% increase in the overall population self-employment rate (12.6% of whom were self-employed before the expansion). These numbers shrink using the local linear method, which is not surprising given the corresponding decrease in estimated SNAP enrollment growth. However, using the narrowest bandwidth the results become insignificant, though the point estimates are positive and the size that would be predicted given the SNAP estimates.

Problems.

As the regression discontinuity bandwidth narrows, point estimates are based on a progressively smaller sample size, which reduces power. For a sufficiently narrow bandwidth, the regression may not have enough power to identify a treatment effect, particularly if this effect is small. This is one potential explanation for the insignificance of the self-employment estimates in Table 3.

Figure 2: Regression Discontinuity (Differenced)

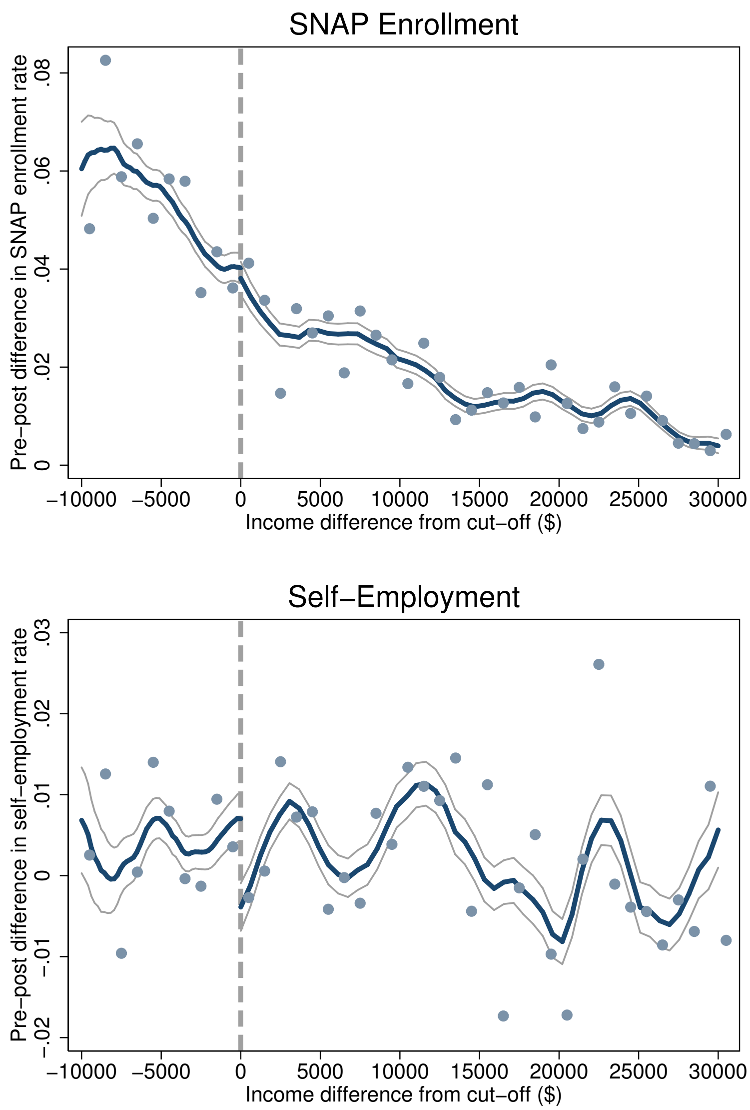

Table 3: Regression Discontinuity Results (Differenced)

Method:	Polynomial	Local linear with bandwidth ...		
		±\$30,000	±\$20,000	±\$10,000
Panel A: SNAP	(1)	(2)	(3)	(4)
$Treat \cdot Post$	0.0483**	0.0395**	0.0353**	0.0232**
	(0.0086)	(0.0077)	(0.0073)	(0.0063)
Pre-policy avg.	0.021	0.021	0.021	0.021
LATE	232.7%	190.3%	170.2%	111.9%
Treated avg.	0.077	0.077	0.077	0.077
TOT	62.4%	51.0%	45.6%	30.0%
Observations	603,058	305,989	236,988	150,499
R-squared	0.045	0.036	0.034	0.026
Panel B: Self-Employment	(5)	(6)	(7)	(8)
$Treat \cdot Post$	0.0184**	0.00858*	0.00708+	0.00392
	(0.004)	(0.004)	(0.0041)	(0.004)
Pre-policy avg.	0.126	0.126	0.126	0.126
LATE	14.6%	6.8%	5.6%	3.1%
Treated avg.	0.067	0.067	0.067	0.067
TOT	27.5%	12.9%	10.6%	5.9%
Observations	603,058	305,989	236,988	150,499
R-squared	0.031	0.008	0.008	0.008

OLS. All specifications include time/state FE and state trends. Data: CPS, 1996-2011.
** 0.01, * 0.05, + 0.1. Robust standard errors clustered at the state level in parentheses.

Another potential problem with this method is motivated by the data. Figure 2 is very noisy, particularly for the self-employment results. In particular, there appears to be a periodic relationship between self-employment differences over time and the assignment rule (income minus the policy cut-off). How much of the estimated treatment effect is being driven by this regular non-linearity in the conditional expectation? Misspecification bias arising from a highly non-linear relationship might push the "treatment effects" up or down, and inference on these results—for example, point estimates or hypothesis testing—might be closely tied to the choice of bandwidth.

In order to address these two concerns, this paper proposes a partitioning strategy that uses information from unaffected portions of the income distribution as counterfactual estimates. This procedure both improves power and controls for misspecification bias, making it more difficult for the econometrician to selectively present bandwidths that confirm desired findings. Section 5 describes this estimator in more detail.

4 Quality and Mechanisms

Business characteristics.

Do this increase in self-employment come from new firm birth or higher survival rates? Columns (1) and (2) of Table 4 separate business-owning households based on whether or not they reported self-employment in the previous year. Existing businesses are actually slightly less likely to survive as a result of the policy, but the effects are not significant and the marginal effects are small. However, non-business-owning households who are newly eligible for the program are 21% more likely to report self-employment after the policy was enacted, though the effects are insignificant, probably because the sample of matched households (for which there is information about the previous year's employment status) is much smaller.

The quality of business ventures is an important topic in entrepreneurship research, and recent evidence suggests there is substantial heterogeneity among entrepreneur quality. Levine and Rubinstein (2012) identify incorporation as an important signal of both a firm's quality (as indicated by future success) and the "seriousness" of the entrepreneur. Since people may enter self-employment as a result of poor labor market prospects and return to wage employment when jobs become available, upticks in entrepreneurship may simply imply a softening labor market. A weak labor market would also push households into public programs like Food Stamps, so these results may conflate the policy's effect with general changes in employment. These brief forays into self-employment, however,

Table 4: Business Source and Quality

	Old firm	New firm					
	Self-Emp.	Self-Emp.	Incorp.	SE % inc.	Workers	Weeks	Hours
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Treat \cdot Post$	-0.0245	0.0046	0.0076*	0.0163	0.148**	1.179**	1.682**
	(0.0304)	(0.0034)	(0.003)	(0.0176)	(0.0485)	(0.345)	(0.254)
Pre-policy avg.	0.69	0.038	0.047	0.40	1.53	47.2	31.1
LATE	-3.5%	12.2%	16.0%	7.2%	9.7%	2.5%	5.4%
Treated avg.	0.58	0.022	0.014	0.46	0.82	43.7	26.9
TOT	-4.3%	20.9%	55.9%	6.2%	18.1%	2.7%	6.2%
Observations	19,600	131,359	603,058	75,558	603,058	528,722	603,058
R-squared	0.024	0.008	0.039	0.094	0.187	0.168	0.089

OLS. Time/state FE and state trends included. RD estimation (differenced), polynomial method (6th-degree function).
Data: CPS, 1996-2011. ** 0.01, * 0.05, + 0.1. Robust standard errors clustered at the state level in parentheses.
Column (3) uses a binary indicator of whether a household owned an incorporated business as the dependent variable.
Column (4) uses the percentage of household income derived from self-employment as the dependent variable.
Column (5) uses the number of employed adults in the household as the dependent variable.
Columns (6) and (7) use the length of the work-year (in weeks) and work-week (in hours) as the dependent variables.

should not be considered “entrepreneurship” in the same sense as more permanent firms.

Are these high-quality ventures? Column (3) of Table 4 uses a binary indicator of whether a household reports having an incorporated business as the outcome variable (these households constitute only 34% of those with businesses and represent 4% of households overall). If the identification strategy is picking up coincidental changes in employment and public assistance use being driven by temporary entrepreneurs, there should be no corresponding change in incorporated ventures. However, column (3) shows a positive and significant coefficient, the marginal effects are half-again as large as those for the general measure of self-employment from Table 3. The number of households with incorporated businesses rose by nearly 0.8 percentage points as a result of the policy, which is a 16% increase from the pre-policy baseline in the total population and a 56% increase among the pre-policy treated group, whose incorporation rates were particularly low. Thus the policy had two important effects on the supply of entrepreneurs: first, it added to the size of their ranks; and second, it shifted the overall distribution toward incorporation, pushing up average firm quality.

Table 3 also presents some results on labor supply. Column (4) uses the share of household income derived from self-employment as a proxy for an entrepreneur's intensive labor supply to the venture. Firms in newly eligible households generated about a 1.6% greater proportion of income for their owners relative to wage income, a 7% increase from before the policy (though this result is not significant). Households also saw large and significant increases in extensive labor supply: treated households saw an 18% increase in the number of employed adults, a 3% increase in the length of the work-year (measured in weeks) and a 6% increase in the length of the work-week (measured in hours).

A simple back-of-the-envelope calculation using the local average treatment effect numbers (to get a conservative and representative figure) indicates the typical household increased their labor supply by 415 man-hours per year as a result of the policy, an increase of 18.5% relative to the pre-policy baseline. Since treated households make up 4.7% of the approximately 115 million households in the US, this adds up to around 2.2 billion man-hours per year, the equivalent of 1.1 million full-time workers.

Mechanisms.

Are the results being driven by a relaxation of credit constraints, a reduction in the risk of leaving wage employment, or some combination of the two? One way to parse these channels is to look at heterogeneity in self-employment growth by uptake of SNAP, since

Table 5: Mechanisms

	Sample restriction:				
	On SNAP	Not on SNAP			
Dependent Variable: Self-Employment	(1)	(2)	(3)	(4)	(5)
$Treat \cdot Post$	0.0051	0.0199**	0.00456	0.00575	-0.0143*
	(0.009)	(0.005)	(-0.0077)	(-0.0047)	(-0.006)
$Income \cdot Treat \cdot Post$			0.00563+		
			(-0.003)		
$Household\ Size \cdot Treat \cdot Post$				0.00383+	
				(-0.0021)	
$(Workers/Children) \cdot Treat \cdot Post$					0.0262**
					(-0.0049)
Observations	14,994	588,064	603,058	603,058	261,470
R-squared	0.017	0.030	0.036	0.037	0.023

OLS. Time/state FE and state trends included. Data: CPS, 1996-2011. Income is in $10,000 increments.
** 0.01, * 0.05, + 0.1. Robust standard errors clustered at the state level in parentheses.
RD estimation (differenced), polynomial method (6th-degree function).

those who are newly eligible but do not enroll receive only a shock to the riskiness of leaving their job, whereas households that enroll experience both a shock to their budget constraint and a reduction in risk. Columns (1) and (2) of Table 4 break down the treatment effect long these lines, and the estimates show that nearly all of the increase in self-employment is being driven by newly-eligible households that do not enroll. Point estimates for enrollees is much smaller than for eligible non-enrollees, and they are not statistically significant. This provides some evidence for the primacy of the risk channel over credit constraints.

Another way to explore these different mechanisms is using household-level variation in financial constraints or budget shocks. Poorer households are the least likely to have access to collateral and therefore most likely to be credit constrained, so marginal effects should be larger for them if credit constraints are binding.[7] Column (3) interacts the $Treat \cdot Post$ variable with income in tens of thousands of dollars, and the estimates are actually positive, meaning marginal effects are larger for richer households (though they are only marginally significant). If credit rationing is important here, than larger shocks to the budget constraint—which represent more severe reductions in credit constraints—should produce larger increases in self-employment. Since benefits are increasing in household size, larger families receive bigger payouts, so column (4) allows the treatment effect to vary by household size.[8] Larger households do in fact experience bigger increases in self-employment, but the estimates are only marginally significant. Finally, households with fewer employed adults per child—that is, a larger dependency ratio—are likely to be more credit constrained and should see larger treatment effects if credit is a binding concern. Column (5) allows the treatment effect to vary by the ratio of workers to children (the inverse of the dependency ratio) and finds a significant, positive result. This indicates households with a *lower* dependency ratio (more workers per child) actually experience the biggest gains in self-employment, casting doubt on the salience of credit constraints. This could suggest families with the greatest potential to self-insure their business using multiple income streams were more likely to enter entrepreneurship as a result of the program, but endogenous fertility complicates this interpretation.

[7] Of course, this elides concerns about absolute risk aversion varying by income level, which unrealistic in a world with CRRA preferences.

[8] In order for this to mean a larger proportional shock to the budget constraint there must be economies of scale in child-rearing, since benefit levels are linear in children.

5 Falsification by Partitioning

One major drawback to RD is the tradeoff between precision of identification and power. Restricting the bandwidth to focus on successively smaller subsets of the data provides a more believable estimate of the treatment effect—these individuals are more plausibly similar to one another than people from opposite ends of the income distribution—but at the cost of throwing away data far from the discontinuity, reducing the sample size and pushing up standard errors. This drawback is reflected in Tables 4 and 5: despite relatively stable point estimates, the results lessen in significance as the bandwidth decreases.

A second problem with RD designs is the possibility of spurious findings arising from the need to approximate a non-parametric conditional expectation with a parametric function. If the conditional expectation is sufficiently non-linear, a local-linear approximation to this function will be biased if the bandwidth is too large. While this bias theoretically approaches zero as the bandwidth narrows, in practice researchers may not know whether the bandwidth chosen is sufficiently small to rule it out as a driver of the results.

Consider a simple case, where some conditional expectation of an outcome variable Y for a running variable X is represented by $g(\cdot)$—that is, E$[Y|X] = g(X)$—and because $g(\cdot)$ is not observed by the econometrician, a local-linear relationship stands in for $g(\cdot)$ in the neighborhood of some discontinuity of interest. Imagine there is some policy such that everyone below a threshold level c is treated and receives outcome η, so that $Y = g(X) + \eta \cdot T_W$ where $T_W = \mathbf{1}[x \leq c]$. Call $\hat{Y}_W$ the linear projection of Y on to X and a dummy for being below the threshold but within the bandwidth W, $\hat{Y}_W = L(Y|T_W, X, x \in W)$, where W is centered around c, and let $\tau(W, g) = \hat{Y}_W - [g(X) + \eta \cdot T_W]$ be the bias of the local linear estimator given a bandwidth and conditional expectation. Clearly $\lim_{||W|| \to 0} \tau(W, g) = 0 \; \forall g$ since the limit of the linear projection is the function itself as the bandwidth narrows. But for a given bandwidth, the RD estimator of η, $\hat{\eta}_{RD} = \lim_{x \to c^+} \hat{Y}_W - \lim_{x \to c^-} \hat{Y}_W = \hat{T}_W = \eta + \tau(W, g)$. What if the econometrician chooses an inappropriately large W because the results confirm some prior or the estimates become too noisy in small bandwidths? Is there a data-driven way to eliminate this bias?

This paper proposes the following falsification check. First, partition the support of $X \backslash W$ into $(||X||/||W||) - 1 \equiv S$ disjoint sets. For each set s of these S sets, create a variable $T_W = \mathbf{1}[x \leq \bar{x}_s]$, where $\bar{x}_s$ is the midpoint of the set. Next, stack each of the sets (renorm $\tilde{x} \equiv x - \bar{x}_s \; \forall x \notin W$). Finally, generate a variable $R = \mathbf{1}[x \in W]$ and project Y onto X, RX, R, T_W, RT_W and a set of dummy variables $\sum_{s \in X} \delta_s$ where $\delta_s = \mathbf{1}[x \in s]$.[9]

The intuition for this procedure is to use the estimated value of T_W for $x \notin W$ as

[9] This allows the intercepts to vary for each subset.

an approximation for the bias $\tau(W,g)$ that would have been observed in a conditional expectation of the form $g(\cdot)$ approximated within a bandwidth of size $||W||$. Clearly $\hat{T}_W|x \notin W \rightarrow \tau(W,g)$ as $||W|| \rightarrow 0$, since both values approach zero. The key assumption of this model is that $\mathsf{E}[\hat{T}_W|x \notin W] = \tau(W,g)$, so that $\hat{T}_W$ can be used to approximate the bias in bandwidths that are not arbitrarily close to zero. If this assumption is satisfied—in particular, if the only thing "special" about W is the presence of a treatment effect η but no other shifts in $g(\cdot)$—then an unbiased estimator of the treatment effect can be obtained using $\hat{\eta}_{PRD} = \hat{RT}_W$, since $\mathsf{E}[\hat{RT}_W] = \mathsf{E}[\eta + \tau(W,g) - (\hat{T}_W|x \notin W)] = \eta$.

Another way to think about this estimator is using a placebo framework. Imagine running placebo RD regressions for regions of the running variable's domain where there is no heterogeneity in treatment (everyone was either treated or untreated). Each of these regressions generates a falsification "treatment effect" that captures misspecification bias arising from the bandwidth choice and non-linearity of the conditional expectation, rather than any real treatment. On average, these values provide an estimate of how biased RD results should be, conditional on bandwidth choice and underlying function. The partitioned RD estimator is the one-step equivalent of running an RD within the bandwidth, then running placebo RDs for unrelated portions of the population, and taking the difference. In this sense, partitioned RD asks how "unusual" the RD estimates in the neighborhood of the policy are given information about the rest of the distribution. This is similar in spirit to searching the distribution of the running variable for the largest significant shift (see Chay and Munshi, 2013) except that the point estimate identifies the difference between the value at the actual threshold and the value elsewhere, rather assigning the entire effect at the jump of greatest significance.

Implementing this process for the context of SNAP enrollment is straightforward. First I demean the assignment rule variable $Rule_{it}$ for each set s where $Rule_{it} \notin W$ by the midpoint of that set, then I define the variable $Real_{it} = \mathbf{1}[Rule_i \in W]$, which indicates being within the neighborhood of the treatment threshold. Then I interact $Real_{it}$ with the variable identifying the treatment effect, which in this case is $Treat_{it}Post_{st}$.

After this partitioning, the partitioned model is

$$\begin{aligned} Y_{ist} = \beta_0 + \beta_1 \cdot Real_{it}Treat_{it}Post_{st} + \beta_2 \cdot Treat_{it}Post_{st} + \beta_3 \cdot Treat_{it} \\ + \beta_4 \cdot Post_{st} + \beta_5 \cdot Rule_i + \beta_6 \cdot Real_{it}Rule_i + B_K + \nu_s + \eta_t + \gamma \cdot t\nu_s + \varepsilon_{ist} \end{aligned} \quad (4)$$

where B_K is a block fixed-effect, and $abs(Rule_{it}) \leq W$ by construction. The coefficient β_1 identifies the treatment effect of the policy net of any misspecification bias arising from the bandwidth choice and non-linearity in the conditional expectation.

Table 6: Regression Discontinuity (Partitioned)

	Bandwidth				
	±\$30,000	±\$20,000	±\$10,000	±\$5,000	±\$5,000
Panel A: SNAP	(1)	(2)	(3)	(4)	(5)
$Real \cdot Treat \cdot Post$	0.0555**	0.0483**	0.0359**	0.0267**	0.0251**
	(0.0088)	(0.0088)	(0.0079)	(0.0061)	(0.0058)
Pre-policy avg.	0.038	0.045	0.055	0.055	0.055
LATE	147.5%	107.1%	65.3%	48.8%	45.8%
Treated avg.	0.077	0.077	0.071	0.055	0.055
TOT	71.7%	62.5%	50.3%	48.4%	45.5%
Covariates	No	No	No	No	Yes
Observations	603,058	603,058	603,058	603,058	600,118
R-squared	0.048	0.049	0.051	0.050	0.082
Panel B: Self-employment	(6)	(7)	(8)	(9)	(10)
$Real \cdot Treat \cdot Post$	0.0206**	0.0180**	0.0141**	0.0139**	0.0091*
	(0.0034)	(0.0039)	(0.0032)	(0.004)	(0.0036)
Pre-policy avg.	0.089	0.083	0.075	0.070	0.070
LATE	23.1%	21.7%	18.9%	19.9%	13.0%
Treated avg.	0.067	0.067	0.064	0.061	0.061
TOT	30.8%	27.0%	22.2%	22.8%	14.8%
Covariates	No	No	No	No	Yes
Observations	603,058	603,058	603,058	603,058	600,118
R-squared	0.031	0.031	0.031	0.032	0.053

OLS. Time/state FE and state trends included. Data: CPS, 1996-2011.
** 0.01, * 0.05, + 0.1. Robust standard errors clustered at the state level in parentheses.

Panel A of Table 5 estimates equation 4 for SNAP enrollment. Using the narrowest bandwidth, which identifies the treatment effect using individuals within \$5,000 of the cut-off and controls for spurious jumps that would have been picked up using a bandwidth of this size, the partitioning strategy finds a 2.7 percentage point increase in SNAP enrollment as a result of the expanded eligibility rules. This is a 48% increase from the pre-policy baseline for both the local average treatment effect and the treatment on the treated; these values are similar because the true "treated" group here are households with income below the cut-off but within \$5,000 of the threshold, and their characteristics are similar to the average household within the bandwidth.

Panel B does the same with self-employment at the outcome variable. The narrowest bandwidth shows a 1.4 percentage point increase in the likelihood of owning a business, an increase of around 20%. The final columns of each panel includes a vector of covariates in the regression. The estimates are very similar in size and significance, with a relatively large increase in R-squared (a near doubling in each case), which indicates they are relevant in explaining variation in the outcome variable yet are orthogonal to treatment status.

6 Model Validation

Covariate balance.

Which model provides the more convincing randomization between treatment and control groups? One simple test of the research design's quasi-experimental quality is to see if the identification strategy picks up shifts in demographic characteristics plausibly unrelated to the policy. Table 6 runs each of the empirical methods discussed so far using a set of covariates as the outcome variables, similar to experimental balance tables presented in randomized control trials.

Column (1) uses difference-in-differences as the identification strategy, with the row name as the outcome variable. Ten out of these eighteen variables (56%) differ significantly at the 5% level in the same way used to identify the treatment effect of the policy. If treatment were randomly assigned, only 5% of these variables should be significant at that level, casting some doubt on the validity of the parallel trends assumption. The differenced regression discontinuity—shown in columns (2) through (4)—improves on this imbalance somewhat, with only 44% of the variables showing up as significantly different. However, partitioned RD—columns (5) and (6)—performs the best in terms of covariate balance, and outperforms differenced RD for comparable bandwidths. For example, using a bandwidth of \$20,000 reduces the number of unbalanced observables by 25% relative to difference-in-RD, and a \$10,000 bandwidth reduces the imbalance by 15%. The improvement is particularly stark relative to difference-in-differences: partitioned RD cuts the number of unbalanced observables in half, suggesting a more convincing pseudo-randomization between treatment and control groups.

Falsification checks.

The typical regression discontinuity framework has a convenient falsification test to determine whether the estimated effect is the result of the policy. Under the assumptions of the model, changes in outcomes for people who are very close to the discontinuity are driven by policy, since within a sufficiently small bandwidth the only differences between individuals is treatment status. A simple test for whether the shift is policy-based is to simulate fake "treatment" using different thresholds. The largest and most significant change in outcomes should be at the true policy threshold, whereas assigning treatment status at non-policy levels should pick up only small and insignificant differences (this is very similar to "mean shift" models; see Chay and Munshi, 2013 for a more detailed example). The procedure should be able to pick out the true policy threshold as the one with the largest significant jump in outcomes. Multiple ranges of significant treatment coefficients at "fake" policies casts doubt on whether regression discontinuity is identifying the real average treatment effect.

Regression discontinuity in differences, however, has a timing dimension since it uses data both before and after the date the policy was enacted. Just as with difference-in-differences, "fake" dates of policy implementation can be used to check whether the estimator is identifying the average treatment effect or a secular year-to-year difference in outcomes between the treated and untreated groups. In difference-in-differences, if the results are being driven by a violation of the parallel trends assumption then assigning post-policy status using a false date should also produce a significant coefficient. However, if the interaction term is significant only when the true policy date is used, the results are more plausibly uncontaminated by coincidental changes between treatment and control groups. This same test can be used in differenced RD estimates using different years as falsification checks.

Because differenced RD has two "dimensions" of falsification—threshold levels and policy implementation years—every combination of false years and levels can be combined to determine the region that is most statistically significant. If the test points are normalized by the true levels and years, a region of large t-statistics in the neighborhood

Table 7: Covariate Balance Across Models

	DID	RD (Differenced)			RD (Partitioned)	
		Polynomial	±$20,000	±$10,000	±$20,000	±$10,000
	(1)	(2)	(3)	(4)	(5)	(6)
SNAP	0.051**	0.048**	0.035**	0.023**	0.0378**	0.0332**
Self-employment	0.010*	0.018**	0.007+	0.004	0.0127**	0.0129**
Age of adults	-3.976**	-4.156**	-3.885**	-2.989*	-2.768*	-2.614*
Elderly	-0.108**	-0.118**	-0.109**	-0.089*	-0.079**	-0.076*
Social Security	-0.089*	-0.100**	-0.090**	-0.067*	-0.069*	-0.067*
Medicare	-0.097*	-0.109**	-0.100**	-0.078*	-0.074*	-0.072*
High school degree	0.037*	0.046*	0.035*	0.030*	0.040*	0.040*
Urban	-0.019*	-0.014+	-0.010*	-0.011*	-0.012	-0.008
Moved	0.020**	0.017*	0.017*	0.015+	0.009	0.009
Number of children	0.092+	0.096+	0.103*	0.075	0.070	0.030
Child support	0.011*	0.010*	0.009+	0.005	0.007+	0.006
Household size	0.097	0.129+	0.117+	0.074	0.100	0.038
Black	-0.006	-0.011	-0.007+	-0.006	-0.012	-0.012
Unemp. benefits	0.014**	0.013*	0.007	0.005	0.009+	0.006
College degree	-0.008	0.013+	0.007	0.003	0.015*	0.010
Disabled	0.001	0.001	0.002	0.003+	-0.001	-0.000
Hispanic	0.018	0.013	0.015	0.015	0.006	0.008
Married	-0.016	0.004	0.006	0.002	0.005	-0.010
Renter	0.030*	0.014	0.023	0.016	0.002	0.005
Veteran's benefits	-0.004	-0.004	-0.004	-0.003	-0.004	-0.004
Unbalanced:						
$\alpha = 5\%$	56%	44%	44%	33%	33%	28%
$\alpha = 10\%$	61%	67%	61%	44%	44%	28%

OLS. Reported coefficients are $Treat \cdot Post$ (DID and RD) or $Real \cdot Treat \cdot Post$ (Partition) with row name as dependent variable. All specifications include time/state FE and state trends. Data: CPS, 1996-2011. ** 0.01, * 0.05, + 0.1 Robust standard errors clustered at the state level are omitted for brevity.

Figure 3: Falsification Tests, SNAP Enrollment

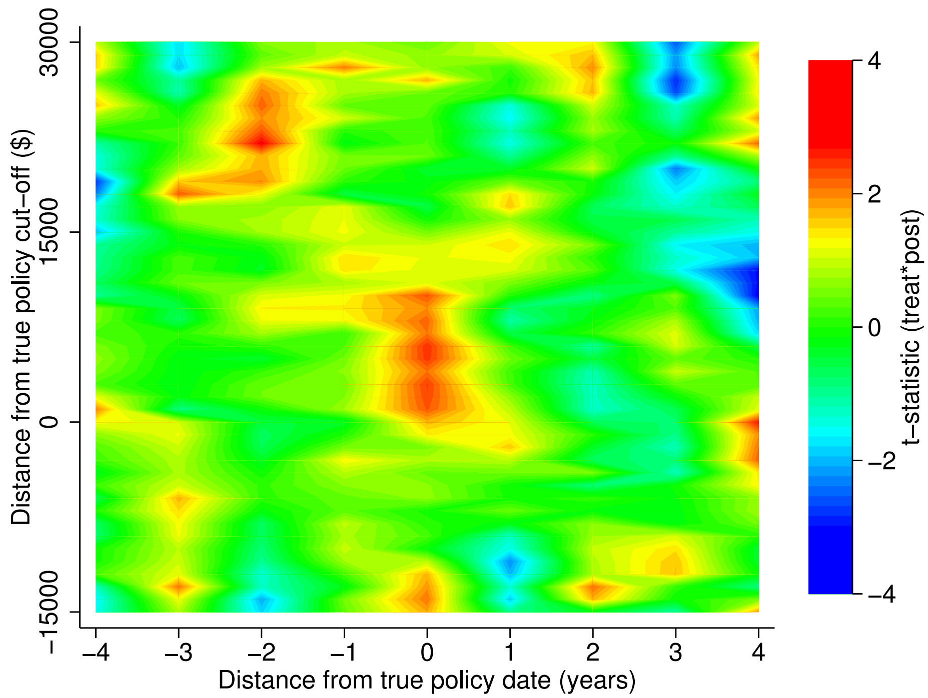

of the origin would provide evidence that the estimator is picking up the true effect of the policy.

A natural way to visualize this three-dimensional relationship is a heat-map. Figure 3 shows the results, with each grid coordinate representing the falsification threshold and year, and the "temperature" of each point representing the t-statistic. These t-statistics come from the interaction term of an estimate of equation (3) using the narrowest bandwidth ($\pm$\$10,000) and with treatment and post-policy status assigned using the information in the grid coordinate, restricting the sample to within one year of the "treatment" date. Redder areas denote a higher positive value and bluer values the opposite. The largest "peak" lies at the actual date the policy took effect, just above the income threshold. The fact that this is not centered at the actual cutoff could be because some eligible households were denied benefits or because income is slightly but consistently misreported in the CPS (Weinberg, 2006).

Monte Carlo Analysis

Understanding how well the partitioning falsification performs relative to RD in more general settings a testing procedure that is not tied to a specific policy, since confounding factors could be driving the results. Monte Carlo analysis using simulated data allows the researcher to control the data-generating process and subject the estimator to a battery of performance tests.

To identify the properties of the partitioning strategy, I use the following Monte Carlo procedure:

1. I draw 10,000 observations of the forcing variable X from the empirical income distribution (taken from the 2010 and 2011 CPS), with the first year of data representing a "pre-policy" period and the second representing the "post-policy" period.

2. Next I construct an outcome $Y = g(X|\xi)$, where I use polynomials of order $k = [1, ..., 7]$ with coefficients $\xi_k \sim N(0,1)$; the natural logarithm; and a sinusoidal function with a period equal to one-fifth of the domain of X for $g(\cdot|\xi)$.

3. I then randomly choose an observation i and assign treatment to all individuals j such that $x_j \leq x_i$ and then subtract x_i from X.

4. I then add a treatment effect η so that $Y = g(X|\xi) + \eta \cdot \mathbf{1}[x_j \leq x_i]$; when modeling Type 1 Error, $\eta = 0$, and for Type 2 Error I use a range of $\eta \in [.05, 1.5]$ where the units are standard deviations of Y.

5. Finally, I use both RD and the partitioned version as the estimator $T(\cdot)$ over a range of bandwidths to generate predictions $\hat{Y}$, in a way similar to equations 3 and 4 but without covariates or post-policy variables.

I then calculate the following statistics for each of the M draws:

- Bias: $\mathrm{E}[(\hat{Y} - Y)/\sigma_y]$, which has the sample analogue $\frac{1}{M}\sum_{m=1}^{M}(\hat{Y}_m - Y)/\sigma_y$;

- Root-mean-squared error (RMSE): $\sqrt{\mathrm{E}\left[\left(\frac{\hat{Y}-Y}{\sigma_y}\right)^2\right]}$, which is calculated as $\sqrt{\frac{1}{M}\sum_{m=1}^{M}\left(\frac{\hat{Y}_m-Y}{\sigma_y}\right)^2}$;

- Type 1 Error: $p\left[\left(\hat{Y}|x_i \leq x_j\right) \neq \left(\hat{Y}|x_i > x_j\right)|T(\cdot), \alpha, \eta = 0\right]$, which has the sample analogue
$$\frac{1}{M}\sum_{m=1}^{M}\mathbf{1}\left[\left(\hat{Y}_m|x_i \leq x_j\right) \neq \left(\hat{Y}_m|x_i > x_j\right)|T(\cdot), \alpha, \eta = 0\right];$$

- Statistical Power: $p\left[\left(\hat{Y}|x_i \leq x_j\right) = \left(\hat{Y}|x_i > x_j\right)|T(\cdot), \alpha, \eta \neq 0\right]$, which is calculated as
$$\frac{1}{M}\sum_{m=1}^{M}\mathbf{1}\left[\left(\hat{Y}_m|x_i \leq x_j\right) = \left(\hat{Y}_m|x_i > x_j\right)|T(\cdot), \alpha, \eta \neq 0\right].$$

Table 8: Monte Carlo, No Treatment Effect

		Bandwidth (in terms of σ_x)						
Panel A: Quadratic $g(\cdot)$		1	0.5	0.25	0.1	0.05	.025	0.01
Type 1 Error:	RD	0.98**	0.96**	0.85**	0.50**	0.40**	0.35**	0.39**
	Partition	0.75**	0.52**	0.29**	0.04+	0.03**	0.02**	0.05
Bias, $sd(y)$:	RD	-0.005**	-0.001**	-0.000**	-0.000**	-0.000**	-0.000**	-0.000**
	Partition	0.166**	0.019**	0.0123	0.011	0.011	0.007	0.011**
RMSE:	RD	0.038	0.007	0.001	0.000	0.000	0.000	0.000
	Partition	0.275	0.046	0.351	0.350	0.350	0.227	0.120
Panel B: Logarithmic $g(\cdot)$								
Type 1 Error:	RD	0.98**	0.97**	0.92**	0.55**	0.43**	0.36**	0.40**
	Partition	0.86**	0.58**	0.33**	0.10**	0.06	0.05	0.04*
Bias, $sd(y)$:	RD	-0.010**	-0.004**	-0.002**	-0.001+	0.000	0.001	0.000
	Partition	-0.050**	-0.019**	-0.003*	0.001	0.004*	0.005*	0.008**
RMSE:	RD	0.046	0.029	0.021	0.011	0.001	0.018	0.001
	Partition	0.094	0.047	0.040	0.047	0.059	0.064	0.098
Panel C: Sinusoidal $g(\cdot)$								
Type 1 Error:	RD	0.98**	0.94**	0.77**	0.54**	0.43**	0.36**	0.40**
	Partition	0.61**	0.61**	0.45**	0.06	0.03**	0.02**	0.04*
Bias, $sd(y)$:	RD	-0.008**	0.000	0.000**	0.000**	0.000**	0.000**	0.000**
	Partition	0.011**	-0.016**	-0.001**	0.000*	0.000	0.000	-0.000
RMSE:	RD	0.053	0.010	0.001	0.000	0.000	0.000	0.000
	Partition	0.066	0.031	0.008	0.002	0.006	0.007	0.007
RD obs. (approx.)		7,120	4,470	2,470	1,030	520	260	110

** 0.01, * 0.05, + 0.1. Robust standard errors calculated using the number of simulations (omitted for brevity).
Tests are for difference from α (Type 1 Error) or zero (bias). $X \sim$ empirical income distribution (2010-2011 CPS).
"RD obs. (approx.)" is observations in RD regressions, averaged over simulations and $g(\cdot)$ functions.
Each cell represents 1,000 draws of 10,000 observations.

If the hypothesis testing procedure is correct, the estimated Type 1 Error should be equal to the size of the test, α (I use a 5% level of significance below). Table 6 shows the results of this procedure when there is no treatment effect ($\eta = 0$) using the empirical income distribution for X.

The results shown in Table 6 display a large reductions in Type 1 Error for the partitioning procedure relative to traditional RD. In fact, for smaller bandwidths, the error approaches the size of the test, $\alpha = 0.05$. While the estimator gains a small amount of bias for larger bandwidths, it is generally on the same order of magnitude as RD.

Robustness tests.

Finally, the empirical estimates can be subjected to more traditional robustness tests. The first test is to allow the conditional expectation to vary across the treatment threshold, which means interacting the assignment rule variable with the treatment dummy. The results of this method for partitioned RD are presented in Table 7 in the appendix, and the results are very similar to those presented above.

Alternatively, this falsification check can be combined with a propensity score matching method, where treatment status is predicted non-parametrically using demographic characteristics of each household. If the partitioning strategy does a poor job of balancing observable characteristics, the results should change noticeably when controlling for the propensity score. The results are shown in Table 7 in the appendix; not surprisingly, the propensity score is a strongly positive predictor of SNAP enrollment and a negative predictor of self-employment, but the treatment effects are still significant are fairly similar to the unconditional estimates from Table 5.

7 Discussion

Entrepreneurs face a number of barriers to entry, including credit constraints and uninsured risk. This paper examines a recent expansion of the Supplemental Nutrition Assistance Program in order to understand which of these constraints is binding. Using three identification strategies to isolate this relationship, I find that becoming eligible for SNAP increased a household's likelihood of enrollment by nearly 3-5 percentage points. Newly-eligible households are 20% more likely to own a business as a result of the policy, driven by an increase in new firm birth of 12%. These tend to be high-quality firms: the marginal effects were particularly strong for incorporated ventures, with the probability of owning an incorporated business increasing by 16% as a result of the policy. The expansion of SNAP also increased the length of the work-year by 2.5% and the work-week by 5% relative to the baseline, a labor supply increase equivalent to 1.1 million full-time workers. I find little evidence that observable characteristics differ between treatment and control groups using the strictest falsification check, and this procedure outperforms RD in Monte Carlo simulations. Finally, I find that the results are driven entirely by newly eligible non-enrollees, suggesting the presence of a large population of would-be entrepreneurs held back by uninsured risk. These findings shed new light on the importance of social programs in shaping labor supply and helping households start businesses.

References

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *The Quarterly Journal of Economics*, February 2004, *119* (1), 249–275.

Besley, Timothy, "Savings, credit and insurance," in Hollis Chenery and T.N. Srinivasan, eds., *Handbook of Development Economics*, Vol. 3 of *Handbook of Development Economics*, Elsevier, 1995, chapter 36, pp. 2123–2207.

Bianchi, Milo and Matteo Bobba, "Liquidity, Risk, and Occupational Choices," *Review of Economic Studies*, 2013, *80* (2), 491–511.

Blanchflower, David G and Andrew J Oswald, "What Makes an Entrepreneur?," *Journal of Labor Economics*, January 1998, *16* (1), 26–60.

Buera, Francisco, "A dynamic model of entrepreneurship with borrowing constraints: theory and evidence," *Annals of Finance*, June 2009, *5* (3), 443–464.

Chay, Kenneth and Kaivan Munshi, "Black Networks After Emancipation: Evidence from Reconstruction and the Great Migration," *Mimeo*, 2013.

Evans, David S and Boyan Jovanovic, "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints," *Journal of Political Economy*, August 1989, *97* (4), 808–27.

Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen, "Entrepreneurial Decisions and Liquidity Constraints," NBER Working Papers 4526, National Bureau of Economic Research, Inc June 1994.

_ , _ , and _ , "Sticking It Out: Entrepreneurial Survival and Liquidity Constraints," *Journal of Political Economy*, February 1994, *102* (1), 53–75.

Hoynes, Hilary W. and Diane Whitmore Schanzenbach, "Consumption Responses to In-Kind Transfers: Evidence from the Introduction of the Food Stamp Program," *American Economic Journal: Applied Economics*, October 2009, *1* (4), 109–39.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Chris Udry, "Agricultural Decisions after Relaxing Credit and Risk Constraints," Working paper 2012.

Levine, Ross and Yona Rubinstein, "Does Entrepreneurship Pay? The Michael Bloombergs, the Hot Dog Vendors, and the Returns to Self-Employment," Working Papers, Haas School of Business, University of California, Berkeley 2012.

McCrary, Justin, "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, February 2008, *142* (2), 698–714.

Olds, Gareth, "Entrepreneurship and Public Health Insurance," *Mimeo*, 2014.

Parker, S.C., *The Economics of Entrepreneurship* The Economics of Entrepreneurship, Cambridge University Press, 2009.

Weinberg, Daniel H, “Income data quality issues in the CPS,” *Monthly Lab. Rev.*, 2006, *129*, 38.

Appendix

Table A2.1: Summary Statistics

Variable	Description	Obs.	Mean	Std. Dev.
SNAP	Receives SNAP benefits	603,058	0.0248633	0.1557085
Self-employment	Has self-employment income	603,058	0.1252914	0.3310493
Gross income	Gross income ($)	603,058	72413.96	66581.78
Black	Has black family member	603,058	0.1276328	0.3336808
Hispanic	Has Hispanic family member	603,058	0.1398986	0.3468821
Household size	Household size	603,058	2.768696	1.445858
Age of adults	Household average age	602,939	46.38724	14.97569
Married	Has married members	603,058	0.6154532	0.4864884
Moved	Moved in the last year	603,058	0.1157839	0.3108406
Urban	Lives in urban area	600,236	0.8195793	0.3845378
Renter	Rents domicile	603,058	0.2637872	0.4406857
High school degree	Has member with high school degree	602,939	0.8898283	0.2689285
College degree	Has member with college degree	602,939	0.3885635	0.4282818
Unemployment insurance	Receives unemployment insurance	603,058	0.0681676	0.2520335
Disabled	Has member receiving disability	603,058	0.0142474	0.1185092
Veteran's benefits	Receives veterans benefits	603,058	0.0224456	0.1481278
Child support	Receives child support	603,058	0.0517844	0.2215917
Social Security	Receives Social Security	603,058	0.2316759	0.4219034
Medicare	Has member receiving Medicare	603,058	0.2178596	0.4127918
Elderly	Has elderly (>65) member	603,058	0.198694	0.3990175
Number of children	Number of children in the household	603,058	0.7772868	1.063063
Workers	Number of working adults	603,058	1.533524	0.9434052
Weeks	Household avg. work-year, in weeks	528,722	46.13667	9.758583
Hours	Household avg. work-week, in hours	603,058	30.99389	7.044649
SE % income	Fraction of income from self-emp.	59,139	0.3960214	0.3511488
Policy year	Year policy was enacted	603,058	2007.038	3.470025
Threshold	Policy threshold (% FPL)	603,058	169.3189	30.70494
Rule	Assignment rule (Inc.$-$Threshold, $)	603,058	46539.69	64957.95
Treat	Treatment status (Rule ≤ 0)	603,058	0.1006918	0.3009205
Post	Post (year $\geq$ postyear)	603,058	0.2933366	0.4552918

Figure A2.1: Density

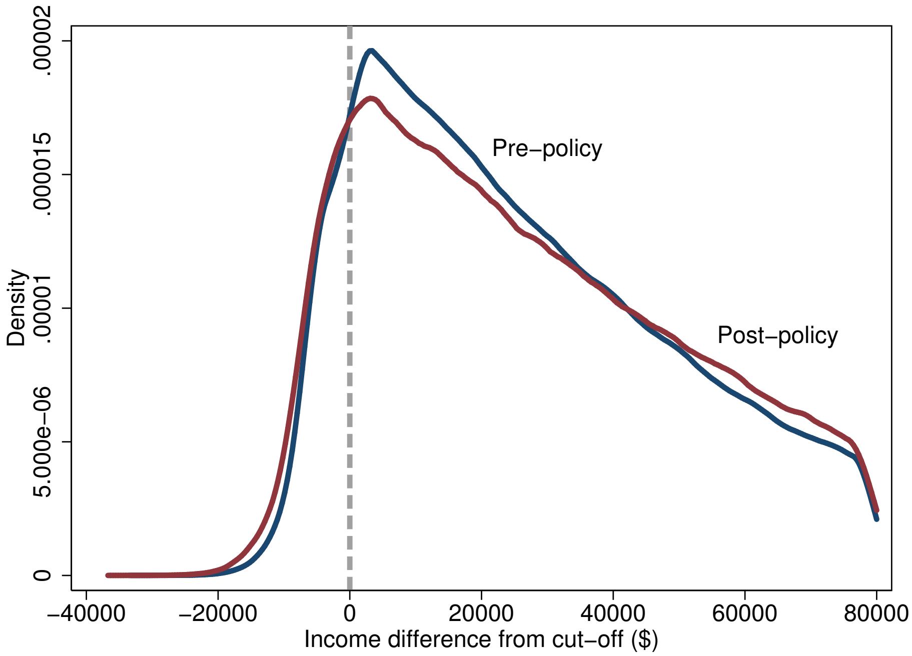

Table A2.2: Partitioned RD and Different Conditional Expectation Slopes

	SNAP Enrollment			Self-Employment		
	±$30,000	±$20,000	±$10,000	±$30,000	±$20,000	±$10,000
	(1)	(2)	(3)	(4)	(5)	(6)
$Treat \cdot Post$	0.0378**	0.0332**	0.0269**	0.0127**	0.0129**	0.0141**
	(0.0092)	(0.0081)	(0.006)	(0.0036)	(0.0032)	(0.0041)
Pre-policy avg.	0.045	0.055	0.055	0.083	0.075	0.070
LATE	83.9%	60.4%	49.3%	15.3%	17.3%	20.3%
Treated avg.	0.077	0.071	0.055	0.067	0.064	0.061
TOT	48.9%	46.6%	48.9%	19%	20.3%	23.2%
Observations	603,058	603,058	603,058	603,058	603,058	603,058
R-squared	0.052	0.052	0.052	0.031	0.031	0.032

OLS. Time/state FE and state trends included. Data: CPS, 1996-2011.
** 0.01, * 0.05, + 0.1. Robust standard errors clustered at the state level in parentheses.
Partitioned RD regressions that allow slope to vary for treated and untreated groups.

Table A2.3: Partitioned RD and Propensity Scores

	SNAP Enrollment			Self-Employment		
	±$30,000	±$20,000	±$10,000	±$30,000	±$20,000	±$10,000
	(1)	(2)	(3)	(4)	(5)	(6)
$Treat \cdot Post$	0.0500**	0.0376**	0.0286**	0.0120**	0.0083*	0.0070+
	(0.0088)	(0.0079)	(0.0061)	(0.0039)	(0.0034)	(0.0039)
Prop. score	0.0767**	0.0798**	0.0796**	-0.273**	-0.271**	-0.271**
	(0.0099)	(0.0098)	(0.0100)	(0.0114)	(0.0112)	(0.0111)
Pre-policy avg.	0.045	0.055	0.055	0.083	0.075	0.07
LATE	111%	68.4%	52.4%	14.5%	11.0%	10.1%
Treated avg.	0.077	0.071	0.055	0.067	0.064	0.061
TOT	64.8%	52.7%	52.0%	18.0%	13.0%	11.6%
Observations	600,118	600,118	600,118	600,118	600,118	600,118
R-squared	0.051	0.053	0.052	0.037	0.037	0.037

Partitioned RD regressions. Time/state FE and state trends included. Data: CPS 1992-2011.
** 0.01, * 0.05, + 0.1. Robust standard errors clustered at the state level in parentheses.