ONLINE APPENDIX

Employed in a SNAP? The Impact of Work Requirements on Program Participation and Labor Supply

Colin Gray, Adam Leive, Elena Prager, Kelsey Pukelis, and Mary Zaki

Online Appendix A: Institutional Details of Virginia SNAP

This Appendix first presents a stylized budget constraint for ABAWDs with and without work requirements, and then discusses exemptions from the work requirements, and reinstatement of the policy in October 2013.

A1. ABAWD Budget Constraint

Work requirements create a notch in the budget constraint as shown in Figure A.1. The graph plots income on the vertical axis against hours not working on the horizontal axis, similar to the representations in Bitler et al. (2006) and Han (2020). For ABAWDs facing work requirements, the budget constraint is represented by ABCFE, with the notch at C. If an ABAWD does not meet the hours threshold, then they lose SNAP after their 3 months of time-limited benefits are used and the budget line drops down to F. Without work requirements, the budget constraint extends from point C to D. There is a slight kink between C and D to reflect the deduction rules that enable SNAP recipients to receive the full SNAP benefit provided their income is below a certain level. This budget constraint abstracts from asset limits (Virginia has limits on liquid assets, such as cash on hand or money in accounts). The graph also does not plot the minimum SNAP benefits, which would change the slope of the budget line close to point B, since our interest is in the notch created by work requirements. The graph draws indifference curves for two hypothetical ABAWDs who do not face work requirements. The person with preferences U^1 is working near the hours threshold without work requirements, and would relocate to point C with work requirements, rather than drop down to the segment EF. This represents the "incentive effect" of the policy. The person with preferences U^2 works only a few hours without work requirements, and would relocate along FE when faced with work requirements given the curvature of their utility function: reaching the hours threshold is too costly for this person and so they exit SNAP and increase their hours. This response constitutes the "income effect" of the policy that removes recipients unable or unwilling to work enough to retain benefits.

A2. ABAWD Work Requirement Exemptions

Individuals are exempt from general work registration if they are younger than 16 years old, 60 years old or older, working 30 hours or more each week, receiving or applying to receive unemployment insurance, serving as a caretaker of a child under the age of 6, temporarily or permanently incapacitated, ill or disabled, regularly participating in an alcohol or substance abuse rehabilitation program, aged 16 or 17 and attending school for at least a half-time basis, aged 16 or 17 but not head of household, enrolled in a recognized school, job skills training, or institution of higher education for at least a half-time basis, already

Figure A.1. : ABAWD Budget Constraint With and Without Work Requirements

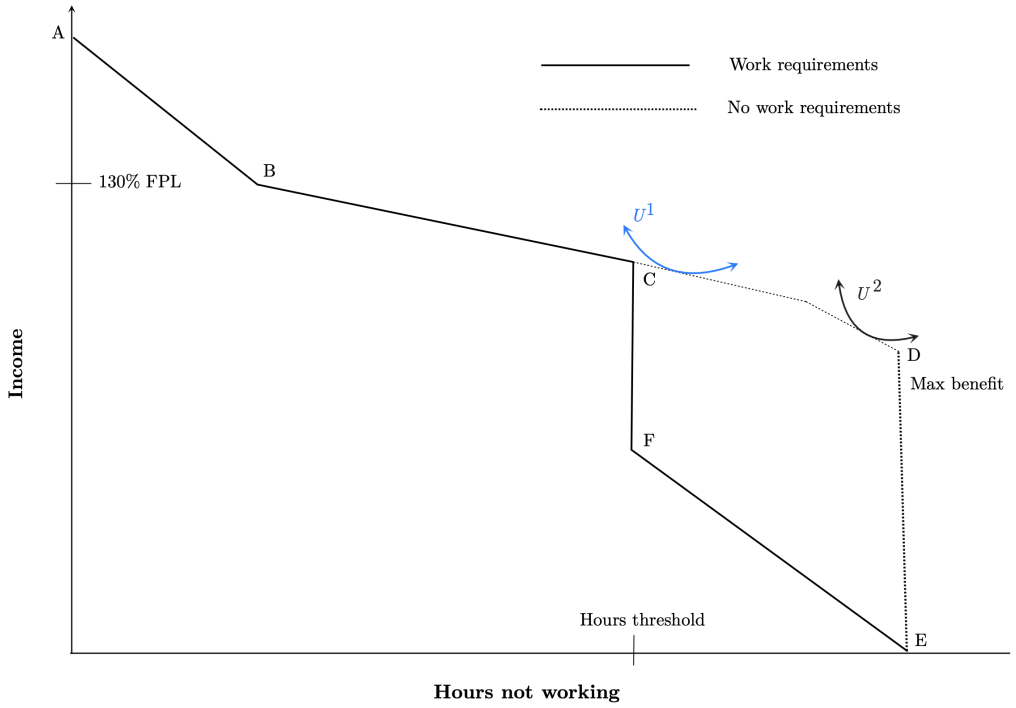

Notes: Figure plots a stylized budget constraint in income-hours space for an ABAWD with work requirements (solid line) and without work requirements (dotted line extension). Without work requirements, the budget set is represented by ABCDE. The number of hours not working is represented on the x-axis, which is a subset of total hours of leisure. Searching for work, which is not part of leisure, does not count towards meeting the work requirements. An ABAWD without work requirements who does not work is represented by the point D, receiving the max SNAP benefit. They can continue to receive the maximum benefit below a certain income threshold due to the deduction rules. Once deductions apply (represented by the slight kink in the dotted line between C and D), SNAP benefits phase out at 30 cents for each dollar of earned income. At point B, the person works enough to earn over 130% FPL and no longer receives SNAP. The budget constraint with work requirements is ABCFE. Under work requirements, the person does not receive any SNAP benefits after their 3 months of time-limited benefits are used unless they reach the hours threshold. Their income in the range below the minimum hours threshold is represented by the segment EF. Working beyond this threshold grants access to SNAP benefits, as represented by the notch in the budget constraint at C. Examples of indifference curves for two types of ABAWDs are drawn tangent to the no-work requirements budget constraint. Under work requirements, ABAWDs with indifference curve U^1 would relocate to C, while ABAWDs with indifference curve U^2 would relocate to E.

complying with another assistance program's work requirements (e.g., TANF or unemployment compensation), or a full-time caretaker of an incapacitated person. Adults are exempt from ABAWD work requirements if they are younger than 18 years old, aged 50 or older, pregnant, medically certified as unable to work, living in a household that includes a child under the age of 18, exempt from general work registration or living in a locality that is exempt from work requirements. Localities (counties and independent cities) may also receive exemptions from the ABAWD work requirements in some circumstances. Specifically, the state office analyzes data and submits a waiver of the requirements for localities that meet qualifications established by the USDA/FNS. A locality may receive an exemption from work requirements if it has a recent 12-month average of unemployment rate above 10 percent, a recent 3-month average unemployment rate above 10 percent, a historical seasonal unemployment rate above 10 percent, a designation as a Labor Surplus Area by the Department of Labor's Employment and Training Administration, a qualification for extended unemployment benefits by the Department of Labor's Unemployment Insurance Service, a low and declining employment-to-population ratio, a lack of jobs in declining occupations or industries, or a recent 24-month average unemployment rate that is 20 percent above the national average for the same 24-month period.

Other than these exemptions, ABAWDs who have already exhausted their allotted SNAP benefits (i.e., 3 months in a 36-month window) can maintain or regain eligibility for SNAP benefits by working at least 20 hours or more per week, participating in an employment services program operated by the Virginia Department of Social Services for 20 hours or more per week (or for at least the number of hours equal to the household's benefits amount divided by the federal minimum wage), participating in an approved work program for 20 hours or more per week, or volunteering for at least the number of hours equal to a household's benefits divided by the federal minimum wage. The state is also annually allotted (by the USDA) a reserve of monthly exemptions based on 15 percent of the number of ABAWD enrollees who live in the state who are not exempted otherwise and do not live in exempted localities. These exemptions may be used by the state to extend the certification period.

A3. Virginia's Reinstatement of Work Requirements

ABAWD work requirements were reinstated in Virginia on October 1, 2013 coinciding with the end of state-wide work requirement exemptions under the American Recovery and Reinvestment Act of 2009 (ARRA), which lasted from April 2009 to September 2013. As shown in Figure A.2, participation rose substantially during the ARRA period but began to fall soon afterwards. The end of ARRA also coincided with an approximately 7 percent drop in the level of SNAP benefits allotted to SNAP recipients in Virginia (Figure A.3). The identification strategy based on regression discontinuity accounts for this benefit

change in estimating the causal effect of work requirements, since the benefit change occurs similarly on both sides of the age 50 cutoff.

Figure A.2. : Monthly SNAP Participation

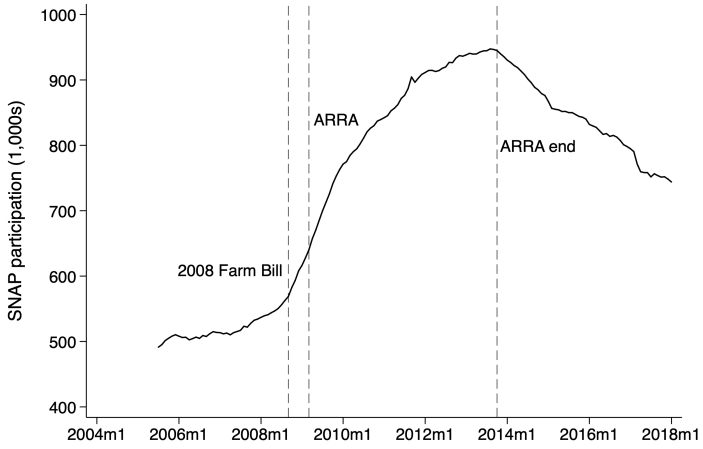

Notes: Figure plots monthly SNAP participation across Virginia measured in thousands. Participation rose substantially during the ARRA period before falling in 2014.

Figure A.3. : Monthly Benefit Amounts, per SNAP recipient

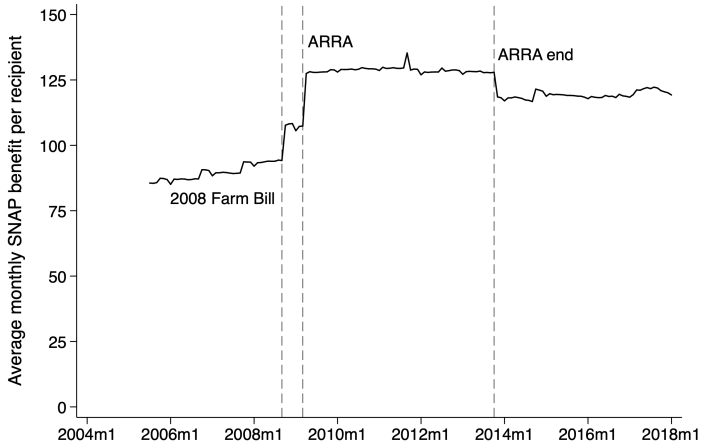

Notes: Figure plots monthly benefits per SNAP recipient across Virginia, calculated as total benefits divided by total SNAP participation. Benefits were reduced by about 7 percent at the end of ARRA.

Prior to the reinstatement of work requirements, individuals typically would

receive 12-month recertification periods. Individuals who began their current benefit receipt period prior to the reinstatement of work requirements continued to receive SNAP benefits up to their recertification date, which would occur after reinstatement. At recertification, their work compliance was evaluated. If they were found to be in compliance of work requirements at recertification, they were given a 12-month recertification period. If they were found not to be in compliance, they were given a 6-month recertification period. This recertification period was composed of an initial partial month of benefits, 3 months allotted for a 36-month window, and a remaining 2 months of exemptions from the reserve of 15 percent exemptions allotted to the state. If after this 6-month period they were again found not to be in compliance with work requirements, they were immediately disenrolled from the program and stopped receiving benefits. The distribution of 15 percent waiver exemptions over time is shown in Figure A.4. These waivers were heavily used throughout most of 2014 but were sparingly used after September of 2014. As a result, Figure A.5 displays a correspondingly large spike in exits of ABAWDs in October of 2014 due to exhaustion of allowable benefit months. If after this 6-month period they were found to be in compliance of work requirements, they were given a 12-month recertification period.

Virginia initially reinstated ABAWD work requirements for all counties in the state on October 1, 2013. However, starting in May 2014, Virginia obtained exemptions for ABAWDs living in 23 counties. Individuals who live in these counties and whose recertifications occurred after the reinstatement of work requirements but before May 2014 received 6 month recertifications. As a result, Figure A.5 shows that ABAWD exits due to exhaustion of allotted benefit months from exempt counties is almost non-existent after October 2014. Those recertifying after May 2014 in exempt counties received 12 month recertifications. In May 2015, Virginia exempted an additional 14 counties.[24]

Finally, ABAWDs who newly enrolled between October 2013 and September 2014 were (theoretically) given 6-month recertification periods. Again, these recertification periods were composed of an initial partial month of benefits, 3 months allotted for a 36-month window, and a remaining 2 months of exemptions from the reserve of 15 percent exemptions allotted to the state. Since 15 percent exemption waivers were not used as readily after September of 2014, those who newly enrolled on or after May of 2014 did not receive a full 6 months of benefits if they did not meet work requirements. Those enrolling between July 2014 and September 2017 generally only received 4-month recertification periods. The variation in recertification periods corresponds to drops in SNAP participation among cohorts of recipients who enroll after October 2013, as depicted in Figure D.1: Figure D.1a shows that the RD point

[24]The following counties and cities were exempted in 2014 and 2015: Bristol, Brunswick, Buchanan, Carroll, Charlotte, Danville, Dickenson, Emporia, Franklin City, Galax, Grayson, Greensville, Halifax, Henry, Hopewell, Lee, Luneburg, Martinsville, Mecklenburg, Norton, Page, Patrick, Petersburg, Pittsylvania, Prince Edward, Prince George, Richmond County, Russell, Scott, Smyth, Southampton, Surry, Sussex, Washington, Williamsburg, Wise, Wythe.

Figure A.4. : Number of ABAWD 15 Percent Exemptions Used

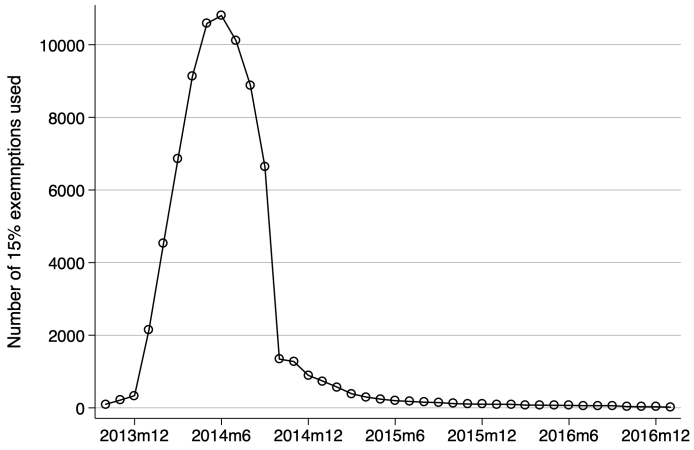

Notes: Figure plots the count of 15 percent waiver exemptions used in Virginia each month. One unit corresponds to one ABAWD being allowed to remain on SNAP for one additional month despite not meeting work requirements. The 15 percent exemptions were primarily used to extend the benefit eligibility of ABAWDs who would otherwise have been removed from SNAP following the reintroduction of work requirements in October 2013.

Figure A.5. : Count of SNAP Exits Due to Failure to Meet ABAWD Work Requirements

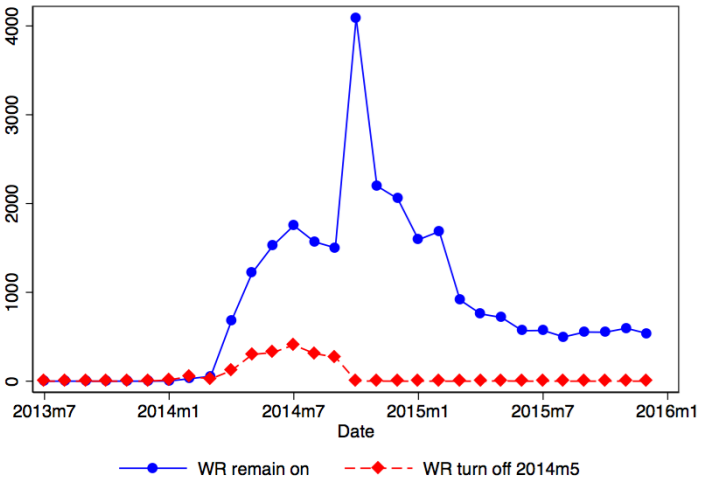

Notes: Figure plots the count of SNAP exits that occur as a result of exhausting allowable benefit months without fulfilling work requirements.

estimate among new participants with 6-month recertifications (October 2013 to April 2014 cohorts) drops substantially in the seventh month. In contrast, Figure D.1b shows that among those with 4-month recertifications (May 2014 to December 2014 cohorts), participation drops in the fifth month. In both cases the RD effects are large and largely comparable in magnitude to the effect estimated using the stock population in the main text.

Online Appendix B: Summary of Related Literature

This Appendix provides further details of the studies described in Section 2.2, and compares them to our paper. We estimate program retention effects that are more than double the magnitudes found in these related papers. Even so, our estimates of labor market effects are comparable to theirs. Our results substantiate a conclusion of truly small labor market effects because of our estimates' tight link to the treatment-on-the-treated combined with accurate administrative data on employment. While some of the differences in results may also arise from across-state heterogeneity in the impacts of work requirements, we view the differences as more likely driven by the data sources used. Table B.1 summarizes the various differences between studies along a number of dimensions.

The public-use versions of the ACS and the CPS aggregate geographic identifiers from low-population counties into larger geographic units, resulting in aggregated units with mixed work requirements policies. The public-use versions also report rounded age, reducing the precision of RD designs using the age 50 cutoff. Stacy et al. (2018) overcome these issues using a restricted-use version of the ACS that reports county-level geography, exact age of respondent on the day of the survey, and the date of survey. However, even in the restricted-use version of the ACS, estimated effects on SNAP participation are diluted because the ACS asks whether a survey respondent was ever on SNAP in the last twelve months, and all ABAWDs (regardless of work status) can receive SNAP benefits for three months when work requirements are in place.

Survey data also severely overstate labor force participation. For example, Harris (2021) reweights survey data to match the demographic composition of ABAWDs in the administrative QC data. Even with this reasonable adjustment, Harris' survey-based ABAWD population has an employment rate of 71 percent, while the ABAWD population in the QC data has an employment rate of 24 percent. Similarly, Han (2020) has a control sample with a 75 percent employment rate. Stacy et al. (2018) have a control sample with a lower—but still high—employment rate of 48 percent from the same survey, potentially because they condition on SNAP participation in the previous year (which may be endogenous). In Ritter (2018), 40 to 58 percent of the CPS-based control sample work more than 20 hours per week while only 10 to 25 percent of the administrative QC-based control sample do. As shown in Table B.1 below, our administrative data on employment is lower and consistent with patterns observed in QC data. For this comparison, we designate as ABAWDs individuals in the QC data who are non-disabled, aged 18 to 49, in childless SNAP units and who are not exempt from work registration.

Table B.1—: Summary Table of Close Contemporaneous Literature

Paper	Labor Market Data Source	SNAP Data Source	Data Structure	States	Identification Strategy
Stacy, Scherpf, and Jo (2018)	ACS (restricted use)	Administrative SNAP data	Labor: Cross-sectional SNAP: Panel	9 States	RD-DD
Harris (2021)	ACS (public use)	ACS (public use)	Cross-sectional	Nationwide	DD & DDD
Han (2020)	ACS (public use)	ACS (public use)	Cross-sectional	Nationwide	DDD
Cuffey, Beatty, and Mykerezi (2021)	CPS (public use)	CPS (public use)	Cross-sectional	Most States	RD
Ritter (2018)	CPS (public use), QC Data (administrative)	N/A	Cross-sectional	29 States	RD
Gray, Leive, Prager, Pukelis, and Zaki [This paper]	Administrative UI data	Administrative SNAP data	Panel	Virginia	Donut RD

Table B.1—: Summary Table of Close Contemporaneous Literature (continued)

Paper	Selected Population for Analysis	Inclusion of Non-treated	Selection Bias? Endogenous variables
Stacy, Scherpf, and Jo (2018)	Ages 25–54 who do not appear to be disabled, have no children under the age of 18, and at or below 250% of the FPL. For labor market outcome analysis, only those who receive SNAP in the 12 months prior to interview date.	Yes	Yes, labor outcomes conditional on SNAP participation, being below 250% of FPL, disability
Harris (2020)	Ages 25–54 who do not appear to be disabled, have no children under the age of 18. Limited to US citizens in continental US who are not institutionalized, in active duty military, or foster care.	Yes	Yes, labor outcomes conditional on disability
Han (2020)	Ages 18–60 who do not appear to be disabled, have no children under the age of 18, and at or below 300% of the FPL.	Yes	Yes, labor outcomes conditional on being below 300% of FPL, disability
Cuffey, Beatty, and Mykerezi (2021)	Ages 20–64, have less than high school education, live in household at or below 250% of the FPL, without children under age 18 and U.S. Citizens.	Yes	Yes, labor outcomes conditional on being below 250% of FPL
Ritter (2018)	CPS: U.S. citizen adults in different age ranges around age 50, who have no children under the age of 18 and who have no High School diploma. QC1: Individuals who do not live with children under the age of 18. QC2: Individuals who do not live with children under the age of 18 and who are not coded as having a disability	CPS - Yes QC1 - Yes QC2 - No	CPS - No. QC1 - Yes, labor outcome conditional on SNAP participation. QC2 - Yes, labor outcome conditional on SNAP participation and disability
Gray, Leive, Prager, Pukelis, and Zaki [This paper]	VA General population (used for case count): SNAP participants who have no children under the age of 18, who are not exempt from work registration and have no known disabilities. Stock population: VA General population members who are enrolled in SNAP in September 2013, who have no children under the age of 18, who are not exempt from work registration and have no known disabilities in that month.	No (at time of sample construction)	No

Table B.1—: Summary Table of Close Contemporaneous Literature (continued)

Paper	**SNAP Participation for control group or at 50 y.o. cutoff**	**Average Labor Outcome for control group or at 50 y.o. cutoff**	**Main Findings**
Stacy, Scherpf, and Jo (2018)	20%	Employment (conditional on SNAP): 48.4%	Work reqt. cause SNAP participation to drop by 3pp but have no significant impact on labor outcomes. No evidence it causes an increase in claims of disability.
Harris (2021)	15.6–18.3%	Employment: 71.5–72.5%	Work reqt. cause SNAP participation to drop by 1.6–1.8pp. They increase employment by 1.1–1.3pp.
Han (2020)	16.60%	Employment: 74.9% Annual earnings: $15,026	Work reqt. exemptions cause SNAP participation to increase by 1.6pp. No significant impact on employment. But find that work reqt. exemptions cause usual weekly hours worked if employed to decrease by 0.388 hours.
Cuffey, Beatty, and Mykerezi (2021)	35%	Employment: 35% Hours worked last week: 12.6 Worked 20+ hours : 32%	13pp increase in employment; increase of 4.7 hrs/week; 13pp increase in working $\geq$20 hours.
Ritter (2018)	CPS - N/A QC - 100%	CPS - work >20hrs: 40–58% QC - work >20hrs: 10–25%	No significant impact on employment except some evidence among the males with disability population from the QC2 sample with 7.3pp effect, significant at the 10% level. Authors find that this result does not stand up to falsification test.
Gray, Leive, Prager, Pukelis, and Zaki [This paper]	Month 0: 100% Month 18: 63.1%	Month 0 employment (SNAP application): 18.6% Month 18 employment (UI): 29% Month 18 annual earnings (UI): $4,276.80	ABAWD SNAP participation decreases by 52% overall and by 38% (24pp) for the stock population 18 months after Work reqt. reinstatement. Homeless are disproportionately impacted. Average employment increases of more than 3.5 percentage points are ruled out. Evidence found of increased earnings near a key eligibility threshold.

Online Appendix C: Additional Tables and Figures

Table C.1—: Descriptive Statistics of SNAP Participant-Months in Full Sample (2005-2016)

	All		Non-ABAWD Adults		ABAWDs	
	Mean	SD	Mean	SD	Mean	SD
Age	24.3	32.8	41.2	16.4	32.5	9.9
Adult	0.55	0.50	1.00	0.00	1.00	0.00
Female	0.54	0.50	0.62	0.48	0.39	0.49
Married	0.13	0.33	0.27	0.44	0.08	0.27
Household Size	3.0	1.7	2.6	1.6	1.2	0.6
Household Head	0.43	0.50	0.76	0.43	0.88	0.32
Homeless	0.02	0.15	0.02	0.13	0.12	0.32
White	0.41	0.49	0.45	0.50	0.37	0.48
Black	0.36	0.48	0.34	0.47	0.37	0.48
Some College+	0.07	0.25	0.12	0.33	0.12	0.32
Has Earned Income (DSS)	0.13	0.34	0.26	0.44	0.16	0.37
Has Unearned Income (DSS)	0.23	0.42	0.34	0.47	0.08	0.27
Avg. Annual Earnings (UI)	4,200	10,463	7,993	13,484	5,959	11,427
Fraction of Months Employed	0.22	0.36	0.39	0.41	0.37	0.38
Ever reported...						
Any Disability	0.15	0.36	0.29	0.45	0.09	0.28
Exempt from Work Registration	0.39	0.49	0.70	0.46	0.26	0.44
Exempt due to Dependent	0.13	0.34	0.25	0.43	0.08	0.27
Medicaid Recipient	0.69	0.46	0.60	0.49	0.29	0.45
TANF Recipient	0.22	0.42	0.16	0.36	0.05	0.21
SNAP E&T Participant	0.04	0.19	0.05	0.21	0.16	0.36
Moved County	0.06	0.24	0.12	0.33	0.03	0.18
N	2,272,827		1,006,065		240,705	

Notes: Table reports descriptive statistics of SNAP participant-months across the whole sample, rather than restricted to the stock population in the main analysis. N denotes count of participant-months. The variables Has Earned Income and Has Unearned Income are reported in DSS files. The variables Avg. Annual Wages and the Fraction of Months Employed are both calculated from UI records.

Figure C.1. : SNAP Participation Survival in Counties with Work Requirement Exemptions

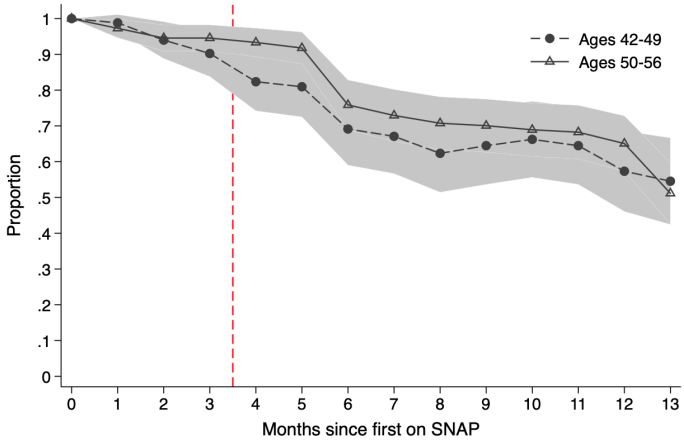

Notes: Figure plots participation survival for ABAWDs aged 42–49 and able-bodied adults without dependents or disabilities aged 50–56 in counties with exemptions for work requirements in May 2014, who have not had a SNAP spell earlier in our sample period, and who first receive benefits between July 2014 and December 2014.

Figure C.2. : SNAP Participation Survival by Work Requirements Status: Participants With Four-Month Recertification Periods

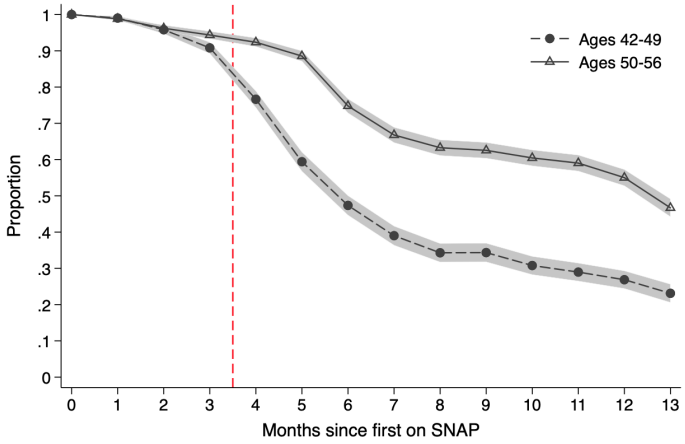

Notes: Figure plots participation survival for ABAWDs aged 42–49 and able-bodied adults without dependents or disabilities aged 50–56 in counties with active work requirements, and who have not had a SNAP spell earlier in our sample period. Work requirements apply to ABAWDs (dashed line), who are required to start meeting them four months after initial entry (dashed red vertical line) in order to continue to receive SNAP benefits. The figure plots participation survival for participants whose SNAP spells begin between July 2014 and December 2014, who are required to start meeting work requirements four months after initial entry. Compare to Figure 1, which plots participation for those required to start meeting work requirements six months after initial entry.

Figure C.3. : RD Estimate of Total SNAP Participation

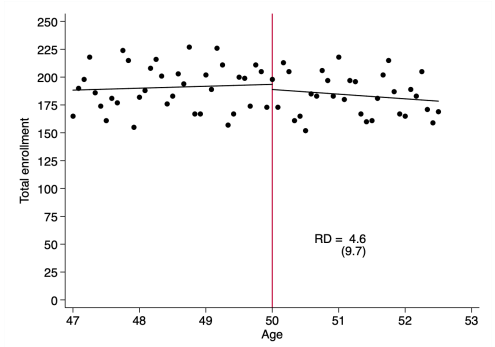

(a) 24 Months Before Work Requirements

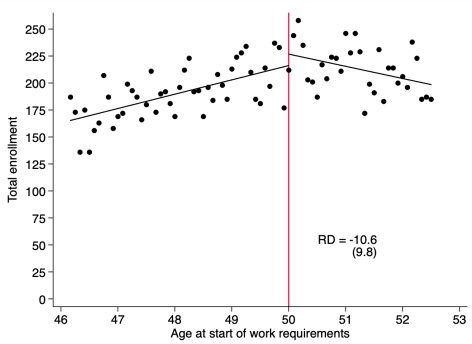

(b) Month of Work Requirements

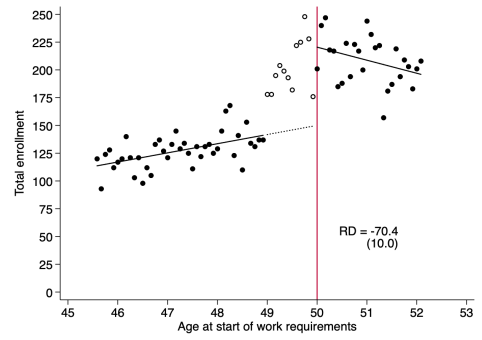

(c) 12 Months After Work Requirements

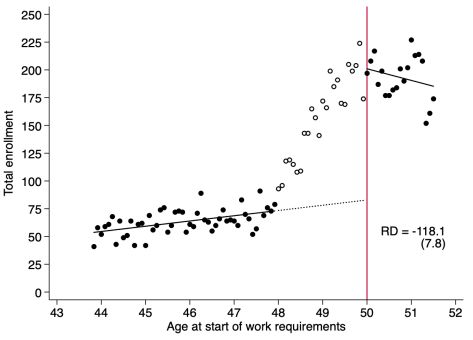

(d) 24 Months After Work Requirements

Notes: Figure visually displays the RD results for total SNAP participation 12 months before, 0 months after, 12 months after, and 24 months after work requirements. In each RD, we define the excluded donut to correspond to those participants whose exposure to work requirements changes between the estimation period and the post-ARRA reintroduction of work requirements. For example, the earlier periods do not require a donut; the period 12 months after the reinstatement of work requirements requires a one-year age donut. Each scatter plot shows total participant counts by age in quarters, and the lines show a linear regression fit on both sides of the eligibility threshold. Standard errors are clustered by monthly age in parentheses. The sample consists of the subset of counties for which work requirements remain on after October 2013.

Figure C.4. : RD Estimates of Total SNAP Participation at Other Intervals

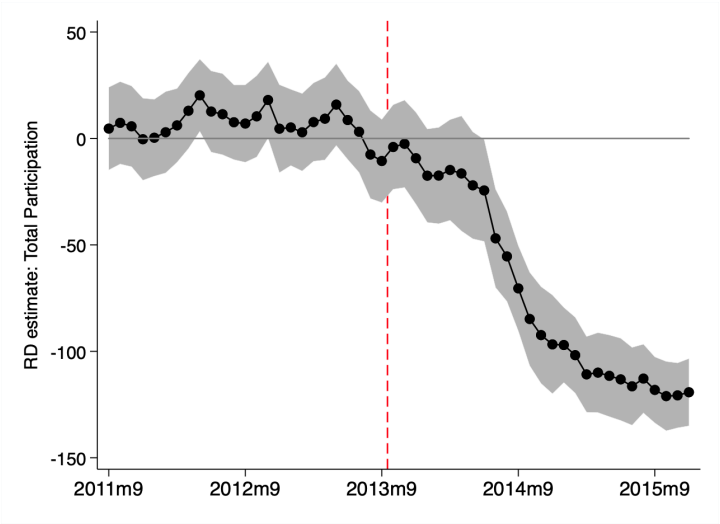

(a) Age Polynomial: Linear

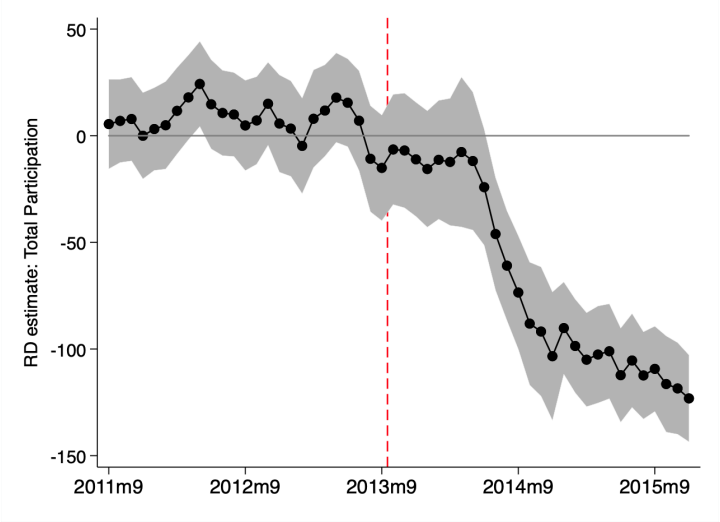

(b) Age Polynomial: Quadratic

Notes: Figure shows RD coefficients for SNAP participation, repeated for other intervals in addition to the baseline interval (18 months after work requirements). In this figure, the coefficient at 2015m3 corresponds to Figure 4a. Panel A presents RD estimates using linear age polynomials and Panel B presents estimates with quadratic age polynomials for robustness. Shaded regions denote 95 percent confidence intervals that cluster standard errors on monthly age.

Figure C.5. : RD Estimates of SNAP Participation at Other Intervals, Stock Sample

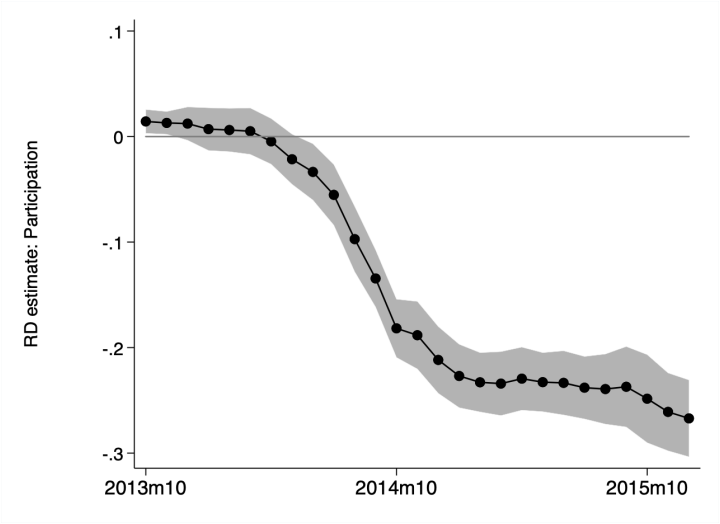

Notes: Figure shows RD coefficients for SNAP participation in the post-ARRA period among the stock population, repeated for other intervals in addition to the baseline interval (18 months after work requirements). In this figure, the coefficient at 2015m3 corresponds to Figure 5a. Shaded regions denote 95 percent confidence intervals that cluster standard errors on monthly age.

Figure C.6. : Robustness to Bandwidth Selection

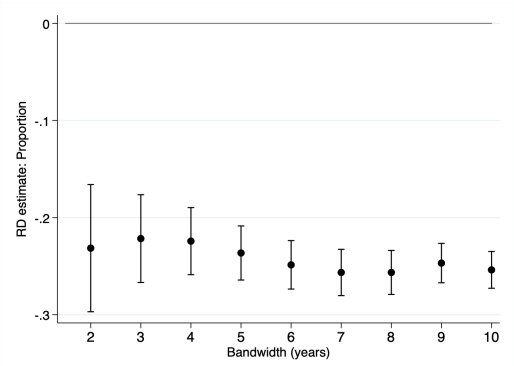

(a) SNAP Participation

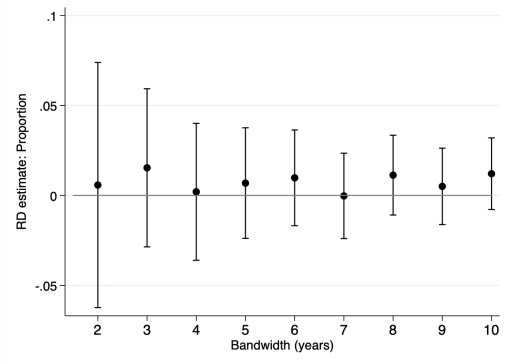

(b) Employment

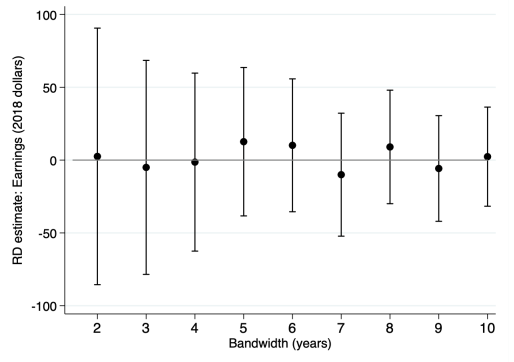

(c) Earnings

Notes: Figures plots the RD estimates 18 months after work requirements were reinstated using different bandwidths. The sample consists of work-registered individuals on SNAP in September 2013 and in the subset of counties where work requirements remain on after October 2013. Earnings are winsorized at the 99 percent level within monthly age.

Figure C.7. : Density of Age at SNAP Enrollment

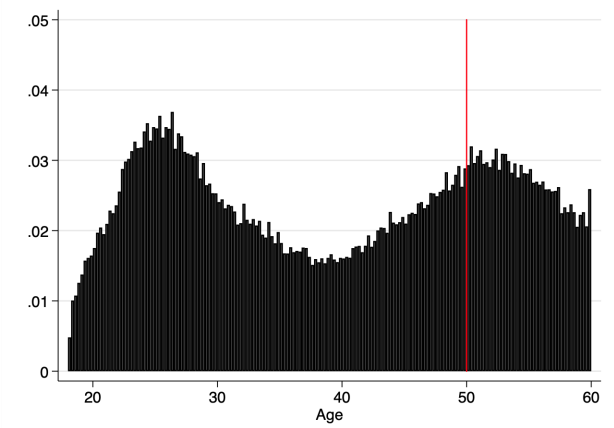

(a) Work Requirements Counties

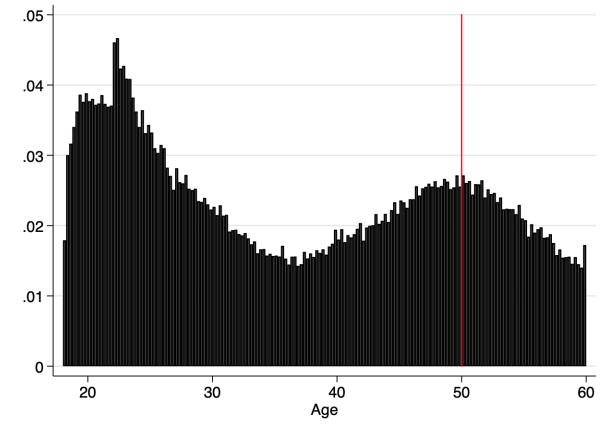

(b) No Work Requirements Counties

Notes: Figures plots the distribution of age at SNAP enrollment within quarterly bins for those in counties with work requirements and those without work requirements. In counties without work requirements, there is no visible discontinuity in the density at age 50. In counties with work requirements, participation appears to be slightly lower just to the left of 50, although the magnitude is small and formal statistical tests (Frandsen 2017) fail to reject the null that the density is smooth at this cutoff. Taken together, there is not strong evidence of selection based on age around the eligibility threshold for work requirements.

Figure C.8. : RD Estimates of SNAP Retention Under Equal Recerticiation Counts

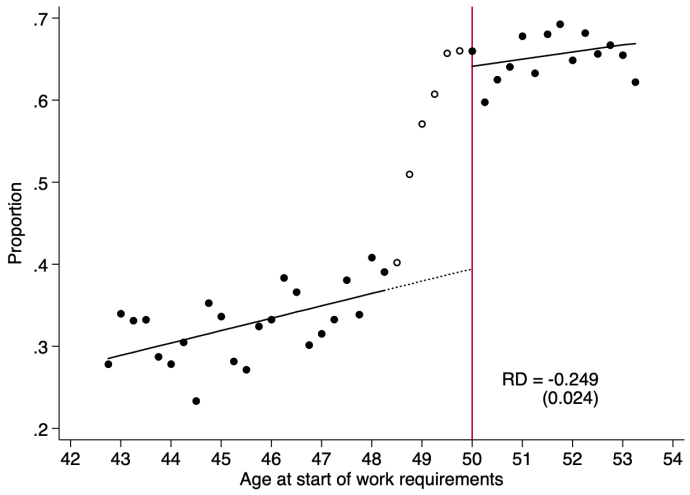

Notes: Figure plots RD results for SNAP participation after eighteen months of work requirements only for the subset of the stock population whose first recertification after the reinstatement of work requirements occurs in the months of October 2013 through March 2014. Participants in this subset of the stock population would have experienced the same number of recertifications (two) if they remain on SNAP 18 months after the reinstatement of work requirements, whether they are in the younger-than-50-group or in the 50-and-older group. Standard errors clustered by monthly age in parentheses. The sample consists of work-registered individuals on SNAP in September 2013 and in the subset of counties for which work requirements remain on after October 2013.

Figure C.9. : RD Estimates of Employment (UI or DSS), 18 Months After Work Requirements

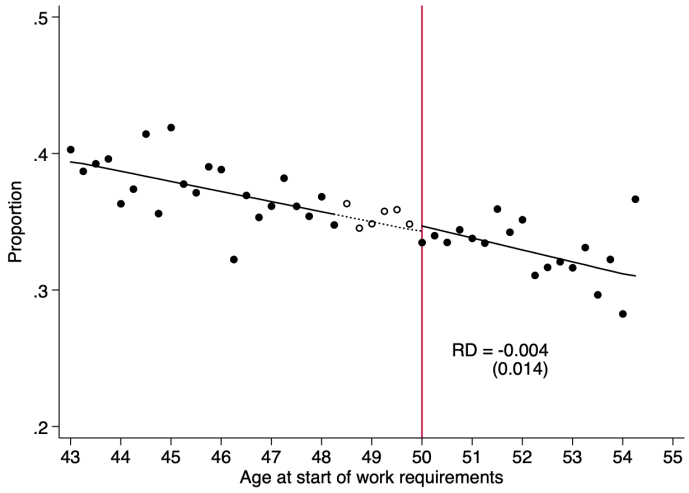

(a) Employment During Work Requirements

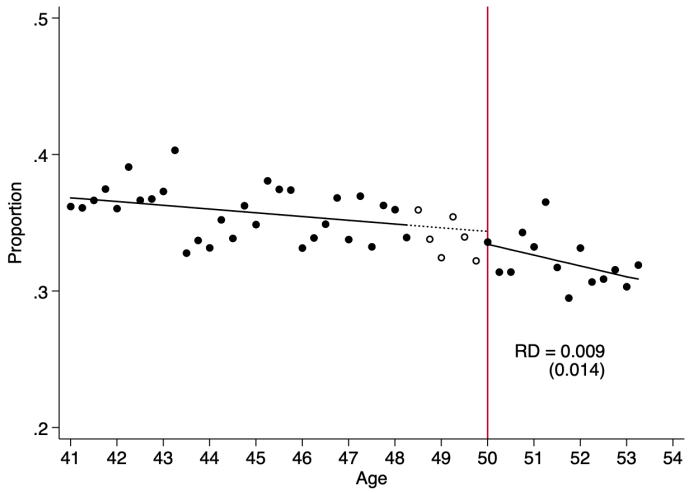

(b) Placebo Test: Employment During ARRA Exemptions

Notes: Panel (a) visually displays the RD results for employment in either the UI data or in DSS-reported earnings after 18 months of work requirements. The scatter plot shows covariate-adjusted means by age in quarters, and the lines show a linear regression fit in months on both sides of the eligibility threshold. Standard errors clustered by monthly age in parentheses. The sample consists of work-registered individuals on SNAP in September 2013 and in the subset of counties where work requirements remain on after October 2013. Panel (b) replicates the same analysis among those participating in September 2011, when the ARRA exemption that suspended work requirements was in effect for an additional two years.

Figure C.10. : RD Estimates of Log Earnings, 18 Months After Work Requirements

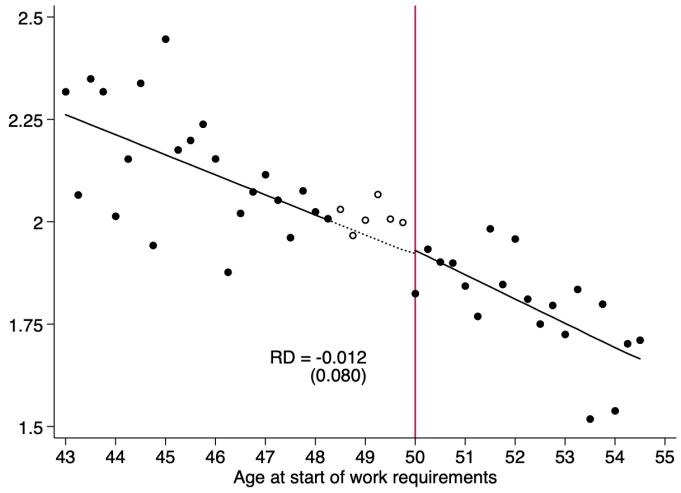

(a) Log Earnings During Work Requirements

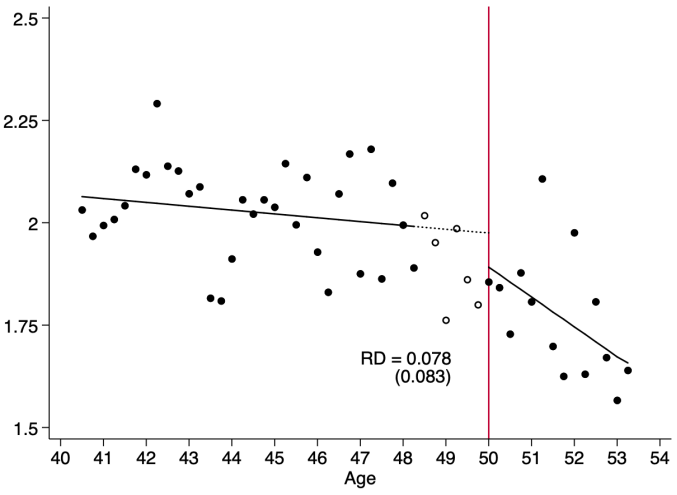

(b) Placebo Test: Log Earnings During ARRA Exemptions

Notes: Panel (a) visually displays the RD results for log earnings after 18 months of work requirements. The scatter plot shows covariate-adjusted means by age in quarters, and the lines show a linear regression fit in months on both sides of the eligibility threshold. Standard errors clustered by monthly age in parentheses. The sample consists of work-registered individuals on SNAP in September 2013 and in the subset of counties where work requirements remain on after October 2013. Panel (b) replicates the same analysis among those participating in September 2011, when the ARRA exemption that suspended work requirements was in effect for an additional two years.

Figure C.11. : RD Estimates of Exempt Status, 18 Months After Work Requirements

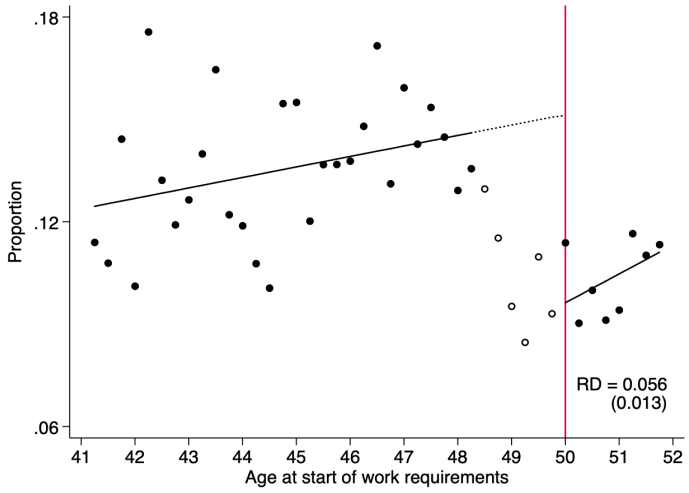

(a) Exemptions During Work Requirements

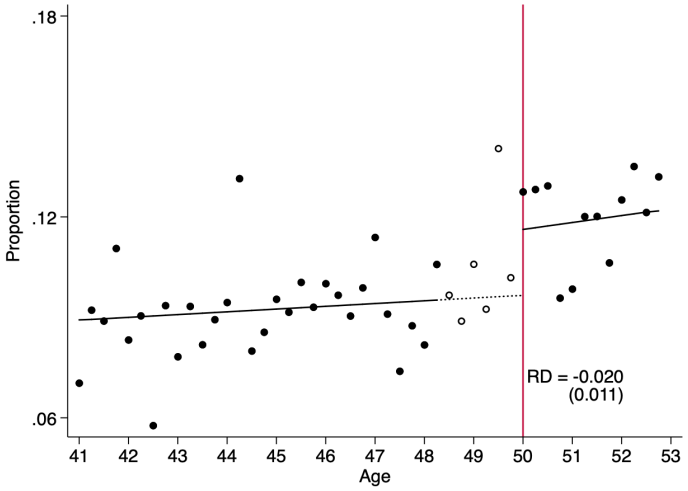

(b) Placebo Test: Exempt Status During ARRA Exemptions

Notes: Panel (a) visually displays the RD results for reported exemptions (except for age) after 18 months of work requirements. The scatter plot shows covariate-adjusted means by age in quarters, and the lines show a linear regression fit in months on both sides of the eligibility threshold. Standard errors clustered by monthly age in parentheses. The sample consists of work-registered individuals on SNAP in September 2013 and in the subset of counties where work requirements remain on after October 2013. Panel (b) replicates the same analysis among those participating in September 2011, when the ARRA exemption that suspended work requirements was in effect for an additional two years.

Figure C.12. : RD Estimates of Employment and Earnings at Other Intervals

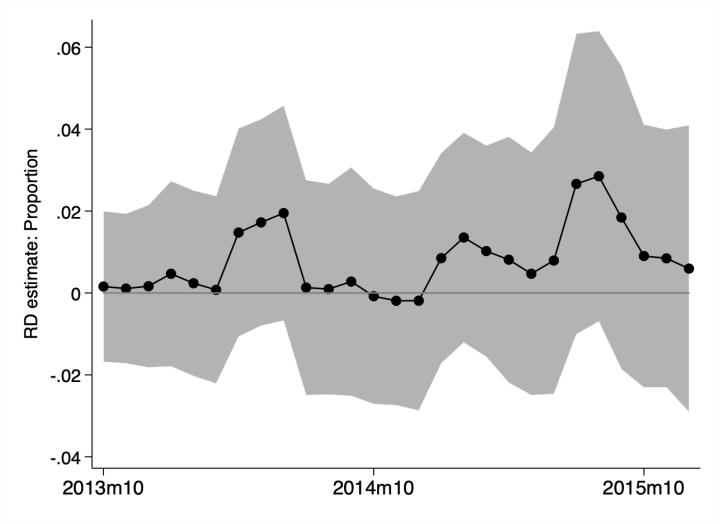

(a) Employment

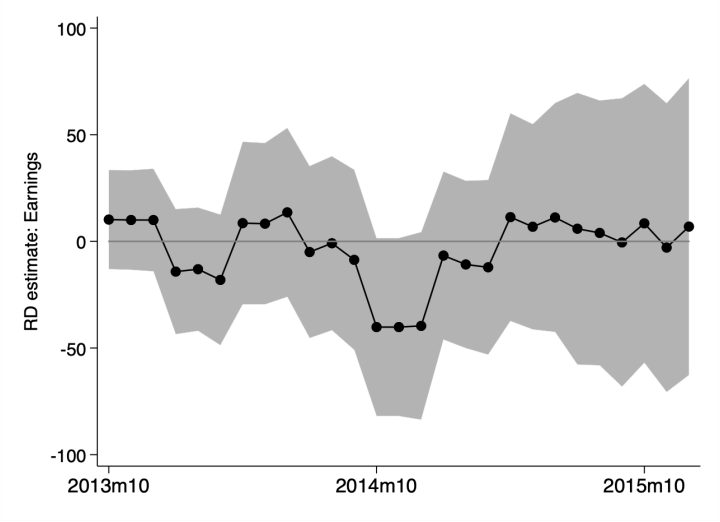

(b) Earnings

Notes: Figure shows RD coefficients for employment and earnings, repeated for other intervals in addition to the baseline interval. Each estimate calculated using a separate MSE-optimal bandwidths on each side of the donut. Shaded regions denote 95 percent confidence intervals that cluster standard errors on monthly age.

Table C.2—: RD Estimates Under Alternative Models

	Linear Uniform		Linear Triangular		Quadratic Uniform		Quadratic Triangular	
	Main	Placebo	Main	Placebo	Main	Placebo	Main	Placebo
	Sept 2013	Sept 2011	Sept 2013	Sept 2011	Sept 2013	Sept 2011	Sept 2013	Sept 2011
Panel A. SNAP Participation								
Discontinuity	-0.234	-0.001	-0.236	0.001	-0.227	0.006	-0.224	0.000
	0.015	0.020	0.014	0.017	0.019	0.024	0.018	0.022
Control Mean	0.632	0.669	0.632	0.669	0.632	0.664	0.630	0.659
N	15,692	13,097	20,144	17,515	24,189	21,270	28,942	24,471
Panel B. Employed								
Discontinuity	0.010	0.007	0.010	0.012	-0.007	0.009	0.007	0.014
	0.013	0.012	0.012	0.011	0.020	0.019	0.018	0.017
Control Mean	0.273	0.277	0.271	0.271	0.263	0.280	0.264	0.274
N	16,840	20,233	19,354	26,079	21,791	23,298	26,383	28,322
Panel C. Earnings								
Discontinuity	-12.2	39.8	-13.6	36.2	-27.7	52.1	-15.7	51.3
	20.7	27.2	17.3	24.8	29.4	34.9	24.8	35.0
Control Mean	365.2	347.7	363.7	340.0	360.1	349.9	362.5	345.8
N	15,930	15,701	22,311	22,108	22,688	24,374	26,592	28,082

Notes: Table shows the main RD estimates under alternative specifications for the kernel and polynomial order. Separate MSE-optimal bandwidths are calculated on each side of the donut for each regression. The first two columns show RD estimates for the stock population (enrolled September 2013) and the placebo stock population (September 2011) using Y_i 18 months later, using the controls described in the text. The third and fourth columns re-weight observations using a triangular kernel. The last four columns replicate this exercise using a quadratic fit on either side of the RD. Control mean is the predicted mean of the corresponding outcome variable immediately to the right of the age 50 threshold (the intercept with the cutoff).

Figure C.13. : RD Estimates of Employment and Earnings at Other Intervals, Without Controls

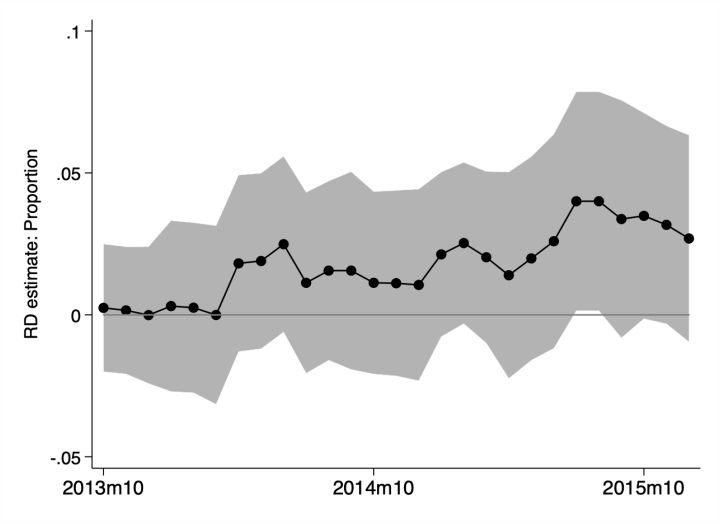

(a) Employment

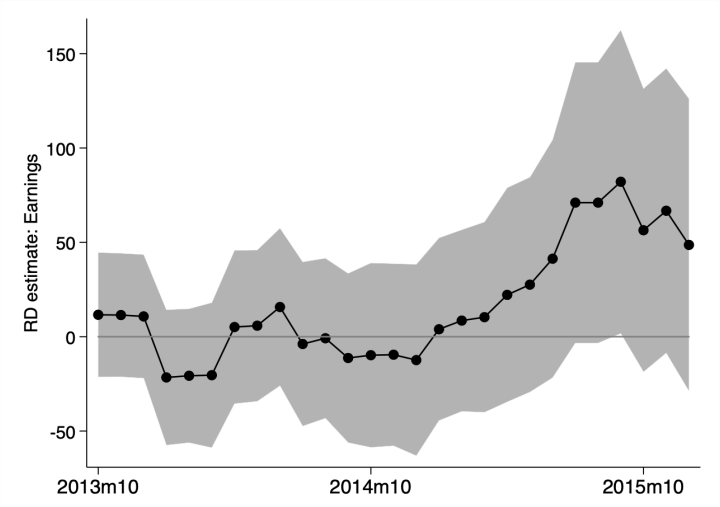

(b) Earnings

Notes: Figure shows RD coefficients for employment and earnings, repeated for other intervals in addition to the baseline interval, in models without controls. Each estimate calculated using a separate MSE-optimal bandwidths on each side of the donut. Shaded regions denote 95 percent confidence intervals that cluster standard errors on monthly age.

Table C.3—: RD Estimates of Key Outcomes Without Controls, 18 Months After Work Requirements

	Main Stock (September 2013)	Placebo Stock (ARRA Period)
Panel A. SNAP Participation		
Discontinuity	-0.245	0.004
	(0.014)	(0.021)
Control Mean	0.629	0.670
N	17,409	13,472
Panel B. Employment		
Discontinuity	0.020	0.010
	(0.015)	(0.015)
Control Mean	0.272	0.277
N	15,570	16,211
Panel C. Employed or Earned Income		
Discontinuity	-0.002	0.014
	(0.019)	(0.018)
Control Mean	0.346	0.333
N	14,705	15,942
Panel D. Earnings		
Discontinuity	10.4	32.3
	(25.4)	(29.6)
Control Mean	368.5	347.1
N	14,943	16,079
Panel E. Log Earnings		
Discontinuity	0.072	0.096
	(0.100)	(0.107)
Control Mean	1.884	1.875
N	16,610	16,079
Panel F. Exemption (Other than Age)		
Discontinuity	0.052	-0.018
	(0.013)	(0.011)
Control Mean	0.095	0.116
N	15,046	19,187

Notes: Table shows regressions coefficients from local linear RD specifications with a uniform kernel, without covariates. Standard errors clustered by monthly age (the running variable) are reported in parentheses. Control mean is the predicted mean of the corresponding outcome variable immediately to the right of the age 50 threshold (the intercept with the cutoff). Employment and earnings are measured from UI records. Log earnings calculated as $\ln(y+1)$. Earnings include those with zero UI earnings, and are winsorized at the 99 percent level by monthly age. The variables Earned Income and Exemption status are reported on DSS records.

Table C.4—: RD Estimates for SNAP Participation and Employment Outcomes at 18 months, by Labor Force Attachment

		Quintile of Predicted Employment Probability in September 2013				
	All	Lowest	2nd	3rd	4th	Highest
Panel A. SNAP Participation						
Discontinuity	-0.245	-0.243	-0.305	-0.192	-0.257	-0.128
	(0.014)	(0.030)	(0.026)	(0.057)	(0.036)	(0.033)
Control Mean	0.629	0.653	0.702	0.622	0.630	0.473
N	17,409	4,329	4,561	2,034	3,512	3,645
Panel B. Employment						
Discontinuity	0.020	-0.014	-0.014	-0.095	0.025	0.031
	(0.015)	(0.022)	(0.024)	(0.072)	(0.032)	(0.026)
Control Mean	0.272	0.127	0.203	0.374	0.300	0.677
N	15,570	4,685	4,389	1,771	2,985	3,852
Panel C. Earnings						
Discontinuity	5.670	-24.518	-89.784	-189.442	17.702	-58.622
	(26.564)	(31.444)	(33.195)	(126.791)	(51.881)	(87.309)
Control Mean	378.536	141.316	259.228	603.446	381.259	1064.881
N	15,595	5,174	5,017	1,852	3,822	3,024

Notes: Table presents the RD coefficient estimates for SNAP participation, employment and earnings for different levels of labor force attachment. Labor force attachment is based on the predicted probability of employment in September 2013, using LASSO regression with data-dependent, theory-driven penalization based on Belloni et al. (2012). The regression includes the following controls to predict employment: indicators for yearly age, indicators for earnings in each month of the previous 7 years prior to Sept. 2013, household size, and indicators for gender, married, private living arrangement, white, black, some college or higher education, reporting earned income on the SNAP application, and reporting unearned income on the SNAP application. Table includes the coefficient, standard error, intercept, and sample size for each specification. Control mean is the predicted mean of the corresponding outcome variable immediately to the right of the age 50 threshold (the intercept with the cutoff).

Table C.5—: Covariate Balance in RD, including 18-month Donut

	Discontinuity	S.E.	Control Mean	% diff	*N*
Female	-0.003	0.023	0.459	-0.7	12,542
White	-0.042	0.018	0.414	-10.1	14,746
Black	0.037	0.021	0.421	8.7	14,490
Married	0.019	0.009	0.096	20.2	20,008
Household Size	-0.004	0.022	1.246	-0.3	16,296
Household Head	-0.011	0.011	0.932	-1.2	13,560
Homeless	0.020	0.011	0.131	15.4	20,213
High School	0.022	0.018	0.538	4.0	18,677
Some College or Higher	-0.014	0.013	0.110	-12.9	17,110
Has Earned Income	0.004	0.011	0.176	2.4	14,943
Has Unearned Income	0.006	0.009	0.091	7.0	16,848
Earned or Unearned Income	0.006	0.015	0.250	2.3	14,567
Fraction of Months Employed, 7yr avg	0.010	0.014	0.351	2.9	14,285
Avg. Annual Earnings, 7yr avg	101.411	255.230	6482.530	1.6	19,321
Fraction of Months Employed, 3yr avg	0.009	0.016	0.281	3.2	14,863
Avg. Annual Earnings, 3yr avg	65.454	316.755	4292.832	1.5	15,960
Number of Months on SNAP	0.254	0.887	27.022	0.9	14,439
Unemployment rate	-0.013	0.029	5.820	-0.2	13,980

Notes: Table presents balance tests of covariates at SNAP enrollment using our "stock" sample. Each row corresponds to a separate regression with that characteristic as the dependent variable. The discontinuity measures the jump in the regression function at age 50. Standard errors are clustered by monthly age (the running variable). Earnings measures are winsorized at the 99th percentile. The Control Mean denotes the mean of that characteristic immediately to the right of age 50. Each regression uses MSE-optimal bandwidths calculated separately for each side of the cutoff and for each outcome, and a uniform kernel to weight observations. Sample sizes vary depending on the bandwidth used.

Figure C.14. : Heterogeneity in RD Estimates of Earnings Using Same Bandwidth, During ARRA Exemptions Placebo Period

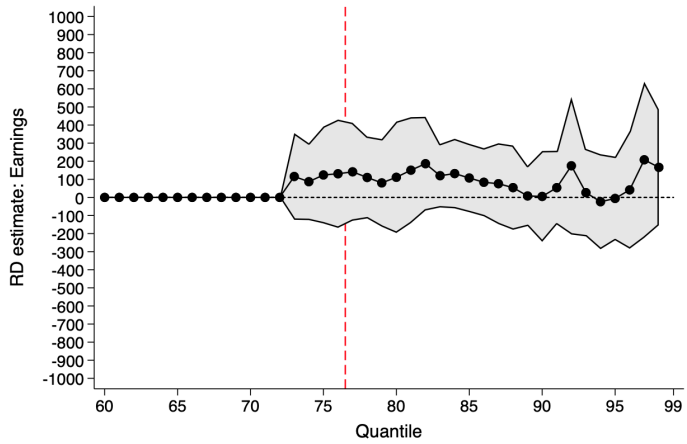

(a) Placebo During ARRA Exemptions: 18-Month Interval

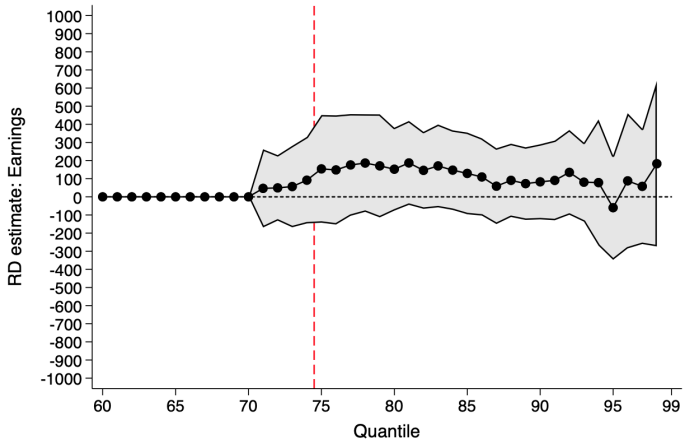

(b) Placebo During ARRA Exemptions: 24-Month Interval

Notes: Figure plots coefficients from individual-level regressions of monthly earnings. Each coefficient is from a separate regression for that quantile using the recentered influence function method of Firpo et al. (2009). Estimates are from the placebo population of individuals on SNAP in September 2011, when the ARRA exemption that suspended work requirements was in effect for an additional two years. Top panel measures earnings in March 2013; bottom panel measures earnings in September 2013. Shading denotes 95 percent confidence intervals using robust standard errors.

Online Appendix D: Additional Analysis Details

D1. Comparison with QC Data

A limitation of this paper is that our estimates may not generalize outside of Virginia. We are not equipped to evaluate across-state heterogeneity in how work requirements are implemented. However, we can compare the compositions of the Virginia and national SNAP populations on observables (measured in the Fiscal Year 2013 QC data). In Fiscal Year 2013, the average SNAP household size in both Virginia and the rest of the country is 1.3. The fraction of individuals whom we would classify as ABAWDs is 7.0 percent in Virginia and 6.3 percent in the rest of the country. The demographic composition is similar except on race. The average age is 36.2 in Virginia and 32.4 elsewhere, and the female fraction is 39.4 percent in Virginia and 40.3 percent elsewhere. The white fraction (47.6 percent in Virginia, 41.6 percent elsewhere) and Black fraction (51.1 percent in Virginia, 36.9 percent elsewhere) are both higher in Virginia than elsewhere, whereas the Hispanic fraction is lower. The fraction with any earned income (unearned income) is 22.6 percent (6.1 percent) in Virginia and 16.4 percent (9.6 percent) elsewhere.

D2. Details of Total Enrollment Decomposition

This Appendix section describes the details of the calculations summarized in Table 2, which decomposes the total enrollment declines documented in Section II.B into three distinct channels:

1) Decreased retention among existing participants.

2) Decreased retention among new enrollees.

3) Deterrence of potential new enrollees.

We perform a series of rough exercises to approximate the relative magnitude of each channel in explaining the overall participation decline. Note that Figure 3b estimates a missing mass of 110.8 beneficiaries per monthly age bin just below age 50. We produce RD estimates for each of the three channels above and compare them to this missing mass. The decomposition into channels need not necessary sum to 100 percent because each channel's contribution is estimated using a separate RD.

To evaluate the first channel, we multiply the number of 50 year olds on the program in September 2013 (226.7) by the main retention effect calculated later in Section III.B (-0.235) and conclude that 53.3 participating individuals per monthly age bin exited due to work requirements. This suggests that the first channel—retention among existing participants—can explain 48 percent (53.3/110.8) of the total enrollment decline.

Figure D.1. : RD Estimates of SNAP Participation in First Year Since Enrollment

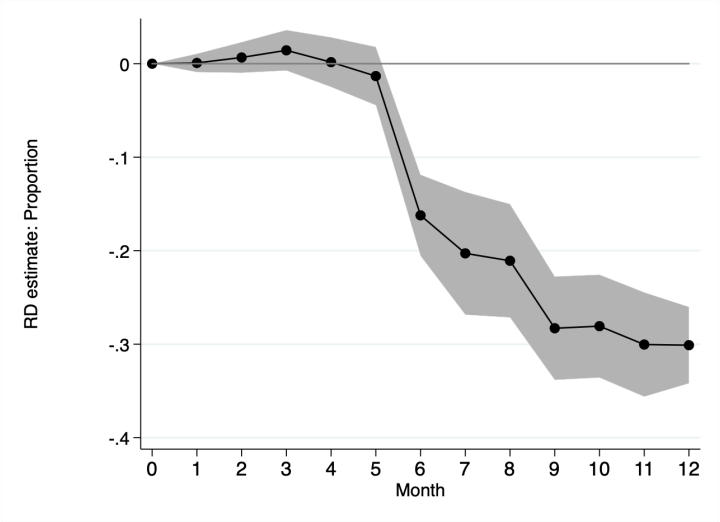

(a) 6-Month Initial Benefit Month Regime

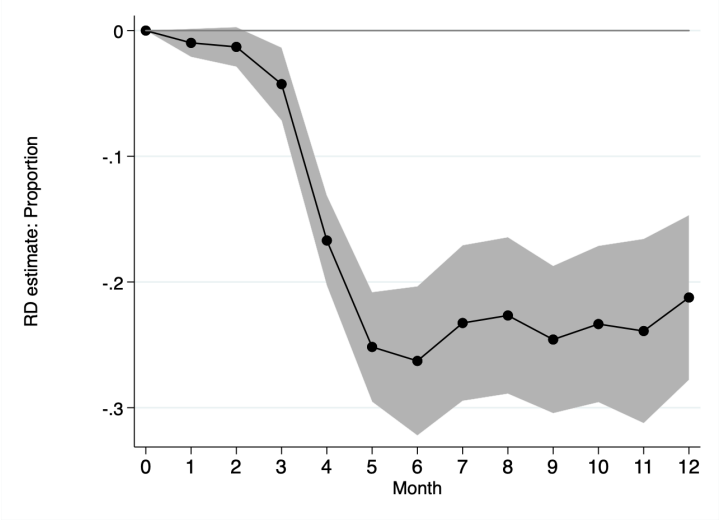

(b) 4-Month Initial Benefit Month Regime

Notes: Figures show RD coefficients for SNAP enrollee cohorts that enter SNAP for the first time since the reinstatement of work requirements in October of 2013, at given points in time since their month of enrollment. Shaded areas represent 95 percent confidence intervals using standard errors clustered by monthly age. Each regression uses the MSE-optimal bandwidth with separate bandwidths calculated on either side of the cutoff.

Figure D.2. : RD Estimates of New SNAP Enrollment by Cohort

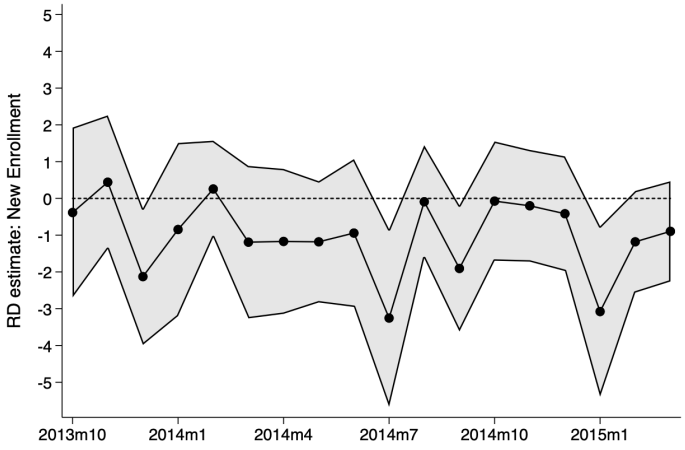

Notes: Figures show coefficients for total new enrollment RDs across successive cohorts of new SNAP entrants. Each regression uses a different MSE-optimal bandwidth, with the bandwidths calculated separately on each side of the cutoff. Shaded regions denote 95 percent confidence intervals that cluster standard errors on monthly age.

The second channel—decreased retention among new enrollees—is difficult to estimate credibly given the possible selection of unobservably different beneficiaries into the program over time. As a very rough approximation, we estimate the loss in retention among new enrollees by estimating separate RDs around age 50 for each subsequent monthly cohort of new entrants after September 2013.[25] Point estimates are shown in Appendix Figure D.1. Multiplying each coefficient by the number of new 50-year-old enrollees in each month yields a sum of −19.8, suggesting a modest role for new beneficiary retention of just below 18 percent (19.8/110.8).

Finally, the third channel—deterrence of potential new enrollees—appears to explain only a small fraction of the total enrollment decline. We estimate (noisy) RDs of total new enrollment in each month from October 2013 through March 2015, and sum the corresponding enrollment drops together. Appendix Figure D.2 shows these coefficients. The coefficients sum to −18, suggesting that new enrollment deterrence can explain less than 17 percent (18.2/110.8) of the total enrollment decline.[26]

These exercises are imperfect but are nonetheless useful. Namely, they provide evidence that deterrence is *not* the primary driver of enrollment declines. Instead, retention of existing and new beneficiaries appear to be the most important channels in total enrollment declines.

[25] We adjust the donut in each RD to exclude those under 50 who turn 50 before March 2015.

[26] We would ideally examine take-up among those *eligible* for SNAP rather than those newly enrolling, but we have no reason to believe this should jump discontinuously at age 50.

D3. Details of Regression Kink Estimation

This Appendix section describes the details of the regression kink (RK) estimates summarized in Section III.B. We estimate RK designs for each of the three main outcomes: participation, employment, and earnings. For each outcome, we begin by estimating a single "joint" RK design using both kinks together. We then follow Card et al. (2016) in estimating two separate RK designs at each policy kink: one at the minimum age when some participants are no longer required to meet work requirements, and one at the maximum age when any participants are required to meet work requirements.

These RK estimates leverage a qualitatively different source of variation from our main donut RD estimates. The donut RD estimates compare participants who had already aged out of work requirements by the time they were reinstated (aged 50 in September 2013) to participants who were just young enough that they would not age out of work requirements until after the 18-month analysis period (aged 48.5 in September 2013). They make no use of the variation provided by participants who were strictly between ages 48.5 and 50 at the time of reinstatement. By contrast, the RK uses the variation "inside" this donut hole by estimating the *change* in the slope of the outcome as participants begin to age out of work requirements (at the left bound of the donut hole) and as they finish aging out of work requirements (at the right bound of the donut hole).

In our setting, the policy rule is deterministic. The share of participants who are required to meet work requirements for at least one month during the 18 months following September 2013 falls from 1 to 0 between ages 48.67 and 49.5 (inclusive). If the share falls linearly, then the RK estimate would be given by the difference in slopes of the outcome with respect to age on either side of a kink, divided by $-1/(49.5 - 48.67) = -1/0.83$. To obtain the difference in slopes of the outcome, we estimate a "joint" RK using slope changes at both kink points. To obtain a single estimate of the slope change, we assume no jumps at the kink points and equal slopes of the outcome variable with respect to age to the left of 48.67 and to the right of 49.5. (We first verify that this is a reasonable simplification. For example, for the participation outcome, the slope to the left is 0.02 (SE = 0.003) and the slope to the right is 0.01 (SE = 0.007).) This yields the following constrained regression equation:

$$y_i = \alpha + \beta(age_i - 48.67) + \gamma \cdot BW_i \cdot (age_i - 48.67) + \delta \cdot BW_i + \varepsilon_i \tag{D1}$$

where BW_i is an indicator for age 48.67 to 49.5, and δ is constrained to be equal to $0.83 \cdot \gamma$.[27]

Table D.1 column 1 reports the estimated slope difference, γ. To check whether

[27] Equivalently, one can estimate $y_i = \alpha + \beta(age_i - 48.67) + \gamma \cdot \max\{0, \min\{0.83, age_i - 48.67\}\} + \varepsilon_i$ without constraints.

these results are consistent with our main donut RD estimates, we focus on the participation outcome, which is the only outcome in which the donut RD does not produce a statistical zero. The estimated slope difference using both kinks is -0.255, making the RK estimate $-0.255 \cdot 0.83 = 0.212$. This is close to the -0.237 estimate obtained using our donut RD specification with a donut between age 48.67 and 49.5 (column 4). For the employment and earnings outcomes, both the "joint" RK estimates and the donut RD estimates are statistical zeros.

Table D.1—: Regression Kink Estimates

	Regression Kink			Regression Discontinuity	
	Both Kinks	Kink at: 48.67	Kink at: 49.5	Donut: 48.67 to 49.5	Donut: 48.5 to 50
Panel A. SNAP Participation					
Estimate	-0.255	-0.480	-0.169	-0.237	-0.245
SE	(0.019)	(0.210)	(0.044)	(0.015)	(0.014)
N	17226	9880	7346	15,762	17,409
Panel B. Employment					
Estimate	0.003	-0.010	0.013	0.009	0.020
SE	(0.016)	(0.133)	(0.038)	(0.015)	(0.015)
N	18953	9767	9186	17,083	15,570
Panel B. Earnings					
Estimate	-1.45	1.64	64.71	0.37	10.36
SE	(25.21)	(146.44)	(49.28)	(20.35)	(25.40)
N	18953	9767	9186	15,954	14,943

Notes: Table shows regression estimates as described in the text. Standard errors clustered by monthly age in parentheses.

In addition to this "joint" RK, we also relax the assumption of equal slopes on either side of the donut hole by estimating separate RK designs at ages 48.67 and 49.5. The motivation for the separate RK designs is primarily institutional: Because of institutional details, the assignment function is only piecewise-linear inside the donut hole, calling into question conclusions drawn from our "joint" RK that assumes a linear assignment function. A change in the length of recertification periods after the reinstatement of work requirements causes a jump in the assignment function between age 48 years 11 months and age 49. From age 48.67 to 48.92, the slope of the assignment function is –1: an additional 1/12 of participants in each age cohort age out of binding work requirements for each month of age. The same is true from age 49 to age 49.5. However, between age 48.92 and age 49, there is a *levels* jump of 0.25 in the fraction of participants who age out of binding work requirements.

We therefore estimate two separate RK designs at each bound of the donut

hole. To obtain the difference in slopes for the first kink, we estimate:

$$y_i = \alpha + \beta(age_i - 48.67) + \gamma \cdot U48.67_i \cdot (age_i - 48.67) + \varepsilon_i \tag{D2}$$

for those who are aged below 48.92 in September 2013. We proceed analogously for the second kink for those aged above 49 in September 2013.

For the RK centered at the lower bound of the donut hole (age 48.67), we use data up to and including age 48.92. This yields an estimated effect of -0.480 for enrollment (Table D.1 column 2). Note that because the slope of the assignment function changes by 1 at each kink, the implied effect size is equal to the raw RK estimate. For the upper bound of the donut hole (age 49.5), we use data starting with age 49. This yields an estimated effect of -0.169 for enrollment. Although they are not statistically distinguishable due to noise, these point estimates are quite different from each other and from our main RD estimates. Taken at face value, they would suggest potentially important conclusions about the effects of *beginning* to age out of work requirements (the estimate from the lower bound of the donut hole) versus *completely* aging out. However, we are cautious in drawing such conclusions. The estimates are noisy and based on small effective sample sizes. The jump in the assignment function at age 48.92 leaves only three age-month cohorts to the right of the kink at age 48.67, and six cohorts to the left of the kink at age 49.5.

Note that we use kink points at ages 48.67 and 49.5 in these RK analyses, rather than our original donut hole bounds of 48.5 and 50 from the main analysis. Briefly, the reason to tighten the donut is that in the RK, it is especially important to place the kinks exactly at the ages where individuals age into and out of being threatened with removal from the program, whereas in the RD we also want to include those who may be affected by other aspects of the reinstatement of work requirements without the actual threat of removal: information about potential removal or shortened recertification periods. The remainder of this Appendix section describes the details.

In our main specifications, we define the donut hole as those who are younger than 50 years old when work requirements are reinstated but who are older than 50 by 18 months after reinstatement. On the age 50 end of the donut, this includes three sets of individuals: (i) those who would be removed from the program during those 18 months if they did not meet work requirements; (ii) those who may be *informed* that they will be removed but who age out of work requirements before the recertification month that would trigger their removal; and (iii) those who receive a shortened recertification period. In the RD, we included the binding work requirements channel, the information channel, and the recertificiation period shortening channel because a priori, there is no reason to assume away the possibility of any channel affecting outcomes. On the age 48.5 end of the donut, we also included individuals who were temporarily subject to binding work requirements, but aged back out of work requirements by the end of

the 18-month analysis period. Again, we include them in the RD because a priori, it is not known whether the impact of work requirements after removal is transient or permanent. Our RD estimates remain similar in magnitude and precisely estimated when we use the narrower RK definition of the donut hole that only includes age cohorts where some (but not all) members would have been *removed* from the program if they did not meet work requirements (Table D.1 columns 4 and 5). This comparison suggests that the removal channel overwhelmingly drives the estimated effects.

In the RK, it is critical to place the kink points at the ages where exposure to the dominant channel changes, because RK estimates rely on highly local changes in slope. Motivated by our finding that the majority of the effect seems to come from the removal channel, we therefore use a narrower age range for defining partial exposure in the RK. We define the donut hole as just those birth cohorts in which some (but not all) members would be required to *meet* work requirements for at least one month prior to turning 50 in order not to be removed from the program. This age interval is 48 years and 8 months to 49 years and 6 months as of September 2013. Whether a particular member of a birth cohort is required to meet work requirements is a function of recertification month and number of "free" SNAP benefit months allotted. (For example, the upper end of the age interval is at 49 years and 6 months rather than 49 years and 11 months because participants were given six "free" months on SNAP before having to meet work requirements.) This definition assumes that the main impacts of work requirements occur during or soon after the first binding month rather than before it (i.e. it assumes no effect of information or shortened recertification period alone); and the marginal impacts of additional binding months are small.

D4. Details of Second Screening Analysis

This Appendix section provides additional details on the second measure of screening summarized in Section III.C. The first measure of screening discussed in the main text assesses whether the exit behavior of participants with certain characteristics is more sensitive to work requirements. Our second screening measure, described in detail here, assesses how the composition of retained participants is affected by work requirements. This measure captures changes due to work requirements in the characteristics of the population of SNAP participants, rather than the differential group-specific exit sensitivities captured by our first measure. Results may differ from those using the first measure when the number of people with a given characteristic is small. In such cases, even large sensitivities may translate to very small compositional changes in the pool of SNAP participants.

To operationalize the compositional measure, we ask which observable characteristics are disproportionately represented among retained participants under work requirements relative to the counterfactual without work requirements. Table D.2 reports estimates from RD regressions on the

proportion of cases with characteristics x among the set of all cases that remain on SNAP in March 2015:

$$x_i = \alpha_1 + \theta_1 U50_i + \gamma_1(age_i - 50) + \gamma_2(age_i - 50) \cdot U50_i + \varepsilon_i \quad \text{(D3)}$$

In this regression, the coefficient of interest is θ_1, which represents the change in composition of retained cases across the age 50 cutoff. Table D.2 shows that, qualitatively similar to the findings in Table 4, work requirements reduce the proportion of individuals with no earned income who remain on SNAP.[28] Also similarly to the findings in Table 4, work requirements cause a greater proportion of those who remain on SNAP to be composed of individuals who have a documented history of having a disability.

Table D.2—: Screening RD by Subgroup, 18 Months After Work Requirements

	Discontinuity	SE	Control Mean	% Diff
Above Median Predicted Earnings	0.026	0.023	0.350	13.8
Female	0.041	0.034	0.472	12.2
Married	0.020	0.018	0.089	24.6
Homeless	-0.032	0.015	0.138	-27.7
White	-0.006	0.024	0.407	0.7
Black	0.002	0.027	0.451	-0.9
Some College+	-0.022	0.020	0.100	-31.7
Has Earned Income	0.041	0.013	0.172	32
Has Unearned Income	0.024	0.015	0.086	25.1
Ever Before UI Recipient	0.001	0.020	0.206	1.3
Ever Before Disability	0.148	0.023	0.179	92.8
Above Median Unemployment Rate	-0.030	0.029	0.418	-7.5
Above Median Previous Time on SNAP	0.001	0.028	0.629	1.2
Above Median Previous SNAP Spell	-0.005	0.026	0.652	-0.1

Notes: Table presents RD estimates of Equation D3. Each row presents results from a separate regression corresponding to the characteristic listed. The first column presents the estimate on the indicator for under 50. Standard errors clustered by monthly age in parentheses are presented in the second column. The third column presents the percentage of 50-year olds who exited SNAP by March 2015 and have the characteristic listed as of September 2013. The last column presents the discontinuity as a percentage of the control mean. The unemployment rate is measured as the county average of the period between October 2013 and March 2015.

[28] Homeless individuals also experience disproportionately large impacts if judged by the individual p-value. However, under this estimate is not statistically significant after adjusting for multiple hypothesis testing using the Bonferroni, Holm, or Benjamini-Hochberg corrections.

D5. Details of the Labor Market Effect Heterogeneity Estimates

This Appendix section provides additional details regarding the estimates of heterogeneous labor market effects in Section IV.B. The estimates use the unconditional quantile regression method of Firpo et al. (2009). Identification in the Firpo et al. (2009) approach relies on the assumption that treatment is exogenous conditional on observables. This is equivalent to the assumptions required for identification in our baseline RD, discussed in Section III.A. If the RD assumptions hold, then exposure to work requirements status is fully determined by observable age and exogenous to other determinants of earnings, and no additional exogeneity assumptions are required for the unconditional quantile regressions.

Note that these estimates do not allow us to identify *which* individuals shifted their behavior as a result of work requirements without stronger assumptions. In other words, it is not possible to say what the counterfactual earnings would be among the people who are at a given quantile in the observed work requirements regime. The identity of the participants at the qth quantile generally will not remain fixed under counterfactual work requirements regimes, except under the assumption that the effect is (weakly) monotonically increasing in the original quantile, which would guarantee rank invariance. The earnings distribution among 49-year-olds stochastically dominates the distribution among 50-year-olds, which is consistent with rank invariance but cannot definitively rule out rank switching.

Figures 7a and 7b suggest that the positive earnings response to work requirements, if it exists at all, takes months to materialize. If real, this delayed response could be explained by three facts. First, it takes time for participants to find (additional) work. Second, improvements in the labor market in mid-2015 may interact with work requirements status. Third, and most relevant, participants can obtain a new 12-month recertification period by temporarily meeting the work requirements after six months. Participants who meet work requirements at the end of their initial 6-month recertification period and then stop working are not removed from SNAP until their next 12-month recertification. In the interim, they can receive several months of benefits without meeting work requirements (see Appendix A). However, a participant who reaches the maximum allowable number of months of not meeting work requirements must subsequently meet them every month to remain on SNAP. Due to this certification schedule, work requirements effectively become more stringent over time (until the work requirements "clock" is reset after 36 months), which is consistent with the observed increase in the earnings impact between Figures 7a and 7b.

D6. Details of the Machine Learning Algorithm

This Appendix section describes how we attempt to attribute the earnings responses documented in Figure 7b to participants who exit SNAP and work more to compensate for the loss of benefits, or to those who work more to retain SNAP eligibility. We refer to these mechanisms as an income effect and an incentive effect, respectively. As discussed in Section IV.B, our RD identification strategy will not yield causal estimates for these mechanisms: examining earnings for those remaining on SNAP in September 2015, for example, involves conditioning on the (endogenous) outcome of not having exited within 24 months of work requirements.[29] Figure 7b provides suggestive evidence of a role for incentive effects, since the largest earnings increases are near the minimum threshold for meeting work requirements, but we cannot rule out strong income effects within this range.

This Appendix attempts a more formal decomposition of income and incentive effects. The intuition is as follows. We attempt to disentangle the mechanisms by classifying participants into three underlying types using machine learning techniques: "never exiters" who would remain on SNAP regardless of whether work requirements are in place, "induced exiters" who would remain on SNAP in the absence of work requirements but exit *due to* work requirements, or "always exiters" who would have exited under either policy regime. RD estimates in the never exiters subsample and induced exiters subsample can then be interpreted as the incentive effect and the income effect, respectively.

The remainder of this Appendix section provides complete details on the machine learning algorithm and the results. The goal of the algorithm is to classify participants into never exiters, induced exiters, or always exiters. The last category primarily consists of participants who would exit SNAP by September 2015 regardless of the presence of work requirements, and who are therefore not useful for estimating the income or incentive effects.

The classification proceeds in two steps. In the first step, we identify and discard the set of participants who would attrit from SNAP by September 2015 even in the absence of work requirements. We identify them by training a LASSO of an indicator for program exit on a wide array of features (listed below) for participants aged 50 to 60, who are not subject to work requirements. We then use these estimates (interpretable as predicted probabilities of exit) to classify participants of all ages into those who would or would not exit in the absence of work requirements. Our main specification selects the classification cutoff to match the empirical probability of exit (39.8%). In the second step, we run an analogous LASSO on participants under 50 (and therefore subject to work requirements) who are *not* predicted to be always exiters, classifying them into

[29]This is an example of the "bad controls" problem, as it involves conditioning on the endogenous outcome of exit. See Section III.C for evidence that beneficiaries who exit due to work requirements differ on observables from those who remain on SNAP.

never exiters or induced exiters. We use these predictions to split both the under-50 and over-50 participants into three groups each, so that RDs within each group have similar composition on either side. We assume away the case of participants who would remain on SNAP if there were work requirements, but would exit SNAP absent work requirements. This assumption is analogous to assuming no defiers in the potential outcomes framework.

More specifically, we begin the process by tuning a LASSO on participants aged 50-60 using ten-fold cross-validation. We implement the algorithm using the *glmnet* package in R using 10-fold cross-validation within each fold to select the tuning parameter λ. We grid search over values of λ between 0.0005 and 0.1 in increments of 0.0005. The features (i.e. covariates) among which the trees select are: indicator variables for female, race is black, race is white, living in a private residence, married, education is less than high school, some high school, high school graduate, some college (omitted category is college graduate), has earned income, has unearned income, ever before had a disability; other variables are household size, 6-month recertification cohort; in the pre-period: fraction of months with wages, sum of pre-period wages, number of months on SNAP; number of months on SNAP in the last 36 months, number of months on SNAP in the last 12 months; wage history from 2005m1–2013m8; and county indicators.

Using this tuned LASSO, we use five folds to obtain predicted values: we run the tuned LASSO on 80 percent of the sample and use the resulting covariates to predict values for the remaining 20 percent.[30] After five iterations of LASSO (with potentially differing covariates), we have a single predicted value for each participant over age 50. We then divide this sample into always exiters and others by selecting a cutoff in the fitted value to match our empirical distribution. We then classify individual *under* age 50 as never exiters or other by taking the average of the five fitted values we obtained from the aforementioned LASSOs.

For the second step, the training sample includes the under 50 individuals that are *not* classified as always exiters. We use an analogous LASSO prediction exercise to further split that sample into induced exiters and never exiters. We again take averages of the five predicted values for each participant over age 50 who is *not* an always exiter, to further classify them as induced exiters or never exiters.

The result of this two-step classification procedure is a sample classified into never-exiters, induced exiters, or neither. We estimate RDs among the never-exiters and among the induced exiters to measure income and incentive effects. The results are presented in the table below.

[30] This helps to prevent bias due to endogenous stratification. For example, suppose we predict earnings and then estimate heterogeneous effect of work requirements on earnings along the range of predicted earnings. Abadie et al. (2013) illustrate how overfitting in the predictive model causes systematic bias in the estimates at low and high values of predicted earnings. In our case, the classification exercise is on participation, which is highly correlated with earnings and therefore could still make us vulnerable to this form of bias. We use this sample-splitting technique to protect our estimates from this systematic bias.

While we selected LASSO as our main specification, we also tried using boosted trees from the R package *xgboost*. The predictive power of the tree was not meaningfully better, and the results were not notably more stable. We therefore opted for the less computationally expensive LASSO.

Table D.3—: Income vs. Incentive Effect at 24 months

	Never exiters (incentive effect)	Induced exiters (income effect)
Panel A. Employment		
Discontinuity	0.034	0.036
95% CI	[-0.063 ,0.125]	[-0.105 ,0.215]
Control Mean	0.199	0.209
Panel B. Earnings		
Discontinuity	41.1	40.8
95% CI	[-113 ,185]	[-163 ,298]
Control Mean	228.0	240.8

Notes: Table shows regression estimates on sub-samples produced by machine learning procedure to decompose earnings changes at 24 months into income effects and incentive effects. "Induced exiters" correspond to SNAP recipients predicted to exit in the presence of work requirements, but not in their absence. "Never exiters" correspond to SNAP recipients predicted to remain on SNAP in the presence of work requirements. Changes in labor market outcomes for these two groups estimate income and incentive effects, respectively. Control mean is the predicted mean of the corresponding outcome variable immediately to the right of the age 50 threshold (the intercept with the cutoff). For each outcome, the 95% CI is calculated by bootstrapping using 1,000 replications, taking the 2.5th and 97.5th quantiles. Estimates of the discontinuity and control mean are calculated by averaging the estimates from all bootstrap replications. Employment and earnings are measured from UI records. Earnings include those with zero UI earnings, and are winsorized at the 99 percent level by yearly age within each calendar month.

Unfortunately, the results are inconclusive. The classification process substantially decreases our sample size and even sophisticated tree methods have limited predictive power. Appendix Table D.3 reports the RD estimates of labor market outcomes within the never exiters (column 1) and induced exiters (column 2). Although some of the point estimates are large, the estimates for both employment and earnings are statistically indistinguishable from zero in both groups. We therefore conclude that work requirements may appreciably increase earnings along a narrow range of the earnings distribution, but the primary mechanism for the earnings increase remains uncertain.

D7. Details of the MVPF Welfare Calculation

This Appendix section provides the details of the welfare calculation reported in Section IV.C. The MVPF for the elimination of work requirements is given by:

$$MVPF = \frac{WTP}{\underbrace{C}_{\substack{\text{direct} \\ \text{program cost}}} + \underbrace{FE}_{\substack{\text{fiscal} \\ \text{externality}}}}$$

where the numerator represents the relevant ABAWDs' willingness to pay out of their own income to eliminate work requirements, and the denominator represents the total cost to the government of eliminating work requirements. Section IV.C defines these terms in detail.

To quantify the MVPF for eliminating SNAP work requirements, consider first the government's cost of eliminating the policy. There are two components of the direct cost per beneficiary. The first component is the change in government spending on program benefits, and is equal to the SNAP benefit for ABAWDs of \$189 per month,[31] multiplied by the fraction of the relevant population who participate in SNAP if and only if work requirements are eliminated. For consistency with our earnings outcomes, we convert the \$189 per month to 2018Q1 dollars, or \$201.12 per month. For consistency with the rest of the paper, we define the relevant population as ABAWDs who would still be enrolled in SNAP after eighteen months, leaving us with 63.2 percent of the ABAWDs who are enrolled at month zero (the natural retention rate reported in Table 5 Panel A). The fraction of the relevant population who participate if and only if work requirements are eliminated is then $0.234/0.632 = 0.370$ (where 0.234 is our main RD estimate of work requirements-induced exit in Section III.B). This yields a program spending increase of $(0.234/0.632) \cdot \$201.12 = \$74.47$ per person per month.

The second component of the direct cost is the cost of administering the program. As currently designed, the SNAP program is costly to administer. For example, the USDA estimates that Virginia's total administration cost in 2016 was \$211 million (Geller et al. 2019). The bulk of these costs, approximately 75 percent, are certification costs.[32] Certification costs fall into two categories: initial certification costs for processing applications to start a new participation spell, and recertification costs for processing spell continuations. Discussions with Virginia DSS staff confirm that certification costs would account for the overwhelming majority of administrative cost changes if work requirements were eliminated. Moreover, the staff believe the marginal cost of verifying work requirements to be small, conditional on performing a certification. We therefore use total certification costs to construct an estimate for administrative cost changes.

[31] The maximum SNAP benefit for a single-person household in the months following the reinstatement of work requirements is \$189. See Section I.C for details.

[32] Other costs include outreach, issuance of the physical EBT cards through which benefits are delivered, and reimbursement of participants' qualified transportation expenses.

The elimination of work requirements could raise or reduce the government's administrative costs. On one hand, the elimination of work requirements would raise program retention, and therefore increase the number of participants requiring recertifications. On the other hand, the elimination of work requirements would lengthen recertification periods, reducing the number of recertifications required per participant. It would also reduce program churn (i.e., the exit and and subsequent re-entry of participants), reducing the number of initial certifications. As discussed below, the administrative cost of an initial certification is substantially higher than the cost of a recertification. We therefore separately estimate the additional recertifications and new applications induced by work requirements using our main donut RD strategy. The resulting estimates suggest that the elimination of work requirements increases the number of recertifications per ABAWD by 0.215 over eighteen months (SE = 0.026) and decreases the number of new applications per ABAWD by 0.114 over eighteen months (SE = 0.026).

We then convert the net change in the number of certifications into dollars using USDA's estimates of certification costs. USDA calculates the total cost of certifications in the data underlying Geller et al. (2019), without distinguishing between initial certifications and recertifications. We obtained the underlying data at the state-year level directly from USDA. We separate the certification-related costs reported by USDA in Geller et al. (2019) into the two components using counts of new and ongoing cases from our administrative data. For example, in 2011, the total certification-related costs reported for Virginia are $\$165,316,905$ in 2018Q1 dollars, and our administrative data contain 167,525 newly initialized cases and 391,539 ongoing cases. Using analogous quantities from the years 2009–2011,[33] this back-of-the-envelope calculation yields estimated costs of \$672.27 per new application and \$153.73 per recertification in 2018Q1 dollars. Discussions Virginia DSS staff confirm that certification of a new application is substantially more resource-intensive than recertification. On net, the elimination of work requirements would reduce administrative costs by $\$672.27 \cdot (0.114/18) + \$153.73 \cdot (-0.215/18) = \$4.26 - \$1.84 = \$2.42$ per person per month. Adding this to the program spending increase yields a total direct cost to the government of $\$74.47 - \$2.42 = \$72.05$ per person per month.

The fiscal externality consists of any changes in net government revenue that result from eliminating work requirements. In our context, a key component is the loss of income tax revenue due to labor market effects. The lost income taxes are given by our estimates of the earnings effect at different durations following the reintroduction of work requirements from Section IV.B. We find no earnings effect 18 months after the reinstatement of work requirements, but a positive earnings effect 24 months after reinstatement. We therefore use the earnings estimate from the 24-month regressions (Figure 7b) to obtain the least favorable MVPF for eliminating work requirements. The most favorable MVPF using the zero earnings

[33]The nationwide suspension of work requirements due to ARRA began in April 2009.

estimate from the 18-month regressions (Figure 7a) is also given at the end of this section. Using the point estimates from the 24-month regressions, the average earnings effect is \$42.24/0.632 = \$66.84 per person per month. If we instead use the upper bound of the 95 percent confidence intervals from the 24-month regressions, we obtain a larger average earnings effect of \$75.08/0.632 = \$118.80 per person per month. This yields the least favorable MVPFs for eliminating work requirements among our calculations, but we view it as less realistic than either the MVPF calculated using the 24-month point estimates or the MVPF calculated using the 18-month estimates. The main text of this section therefore uses the \$66.84 magnitude from the 24-month point estimates, and provides the two more extreme MVPFs at the end.

The government's loss of revenue is \$66.84 multiplied by the tax rate on these earnings. For a single unmarried earner working 80 hours per month at the minimum wage (annual earnings of \$6,960), the average tax rate is approximately 16.5 percent.[34] Thus, the fiscal externality from lost tax revenue due to earnings responses is $0.165 \cdot \$66.84 = \$11.03$ per relevant ABAWD per month. This amount results in a denominator of $(0.234/0.632) \cdot \$201.12 - \$2.42 + \$11.03 = \$83.08$ per relevant ABAWD per month.

The fiscal externality could, in principle, also include changes in per-participant monthly benefits as a result of partial phase-out of benefits with rising earned income. As discussed in Section I.A, benefits are reduced by 30 cents for each additional dollar of income, in addition to a 20 percent earned income deduction. This implies that SNAP benefits effectively decline by 24 cents for each additional dollar of earned income. In the MVPF calculation, we assume that the earnings impacts of work requirements are a result of income effects. This means earnings increase only among work participants who exit as a result of the work requirements, so their income changes are already accounted for in the tax revenue calculation. If the earnings impacts are instead driven by incentive effects, then the government has an additional fiscal externality of $0.24 \cdot \$66.84 = \$16.04$ for a total denominator of $(0.234/0.632) \cdot \$201.12 - \$2.42 + \$11.03 + \$16.04 = \$99.12$ per relevant ABAWD per month.

The numerator of the MVPF is a relevant ABAWD's willingness to pay to eliminate work requirements. In the literature, the numerator is typically equal to the value of the benefits change, as any behavioral response to the policy change is assumed to have zero impact on utility. This assumption of zero impact relies on the envelope theorem combined with benefit changes being "small." In our setting, the policy change being considered does not change the *amount* of benefits received conditional on receipt. Instead, the elimination of work requirements gives benefits to new participants and gives working

[34] To calibrate the tax rate, we note that few non-disabled, non-elderly childless adults are eligible for social programs that would implicitly tax their income. Therefore, average tax rates are primarily composed of payroll taxes of 15.3 percent applied to all earnings, plus a 10 percent income tax applied to earnings above the standard deduction (\$6,300 in 2015).

beneficiaries the option to work less. The former is given by the value of the benefits, multiplied by the fraction of the relevant population that gains benefits if work requirements are eliminated: $(0.234/0.632) \cdot \$201.12 = \$74.47$.[35] The latter benefit is usually ignored in MVPF calculations by assuming that any utility changes from reoptimizing behavior are second-order. In the setting of SNAP work requirements, the reoptimization may lead to substantial utility changes (see Figure A.1). The utility change for this group is bounded between \$0 and \$201.12 per ABAWD.[36] The largest group to whom these bounds could apply constitutes approximately $0.13/0.632 = 21$ percent of the relevant population, where 0.13 is the fraction of percentiles for which we detect earnings impacts in Figure 7b. Thus, the numerator of the MVPF is bounded between \$74.47 and $\$74.47 + 0.21 \cdot \$201.12 = \$116.71$.[37]

We now have approximations for the numerator and the denominator of the MVPF for eliminating work requirements. Using our largest estimate of the earnings response twenty-four months after work requirements (see Figure 7b), the MVPF is bounded below by $\$74.47/\$83.08 = 0.90$ if individuals pay no utility cost of working more as a result of work requirements. If individuals pay a utility cost for the large behavior change of working more, the MVPF is bounded above by $\$116.71/\$83.08 = 1.40$.[38]

If the true earnings response is zero, as we find eighteen months after work requirements (see Figure 7a), then the MVPF is bounded between $\$74.47/\$72.05 = 1.03$ and $(\$74.47 + 0.13 \cdot \$201.12)/\$72.05 = 1.40$. The strongest case against eliminating work requirements—that is, the strongest case in favor of keeping them—would use the upper bound of our 95 percent confidence interval from the 24-month earnings estimates (see Figure 7b). The MVPF

[35]Two caveats are in order. First, this amount is not a "small" change for many beneficiaries, so this transfer is directly proportional to a utility gain only if income effects are small (e.g., quasi-linear utility). Second, this willingness to pay will be lower if ABAWDs value a dollar of SNAP benefits at less than a dollar of income. Our calculation assumes that a dollar of SNAP benefits is valued equally to a dollar of cash. This assumption is consistent with studies find that a dollar of SNAP benefits is spent like a dollar of cash (Hoynes and Schanzenbach 2009, Hoynes et al. 2015). Other work suggests that the marginal propensity to consume food (MCPF) out of SNAP benefits is lower than the MPCF out of cash (Hastings and Shapiro 2018), but this does not imply that SNAP participants value a dollar of benefits at less than a dollar of cash.

[36]We obtain bounds on the effort cost of work requirements-induced work using a revealed preferences argument. The lower bound is equal to the additional income; if the effort cost were less than the additional income, then these participants would work more even in the absence of work requirements. The upper bound is equal to the sum of the additional income and the value of SNAP benefits, which is what participants stand to gain from working when there are work requirements. At the lower bound, the utility cost of working exactly offsets the income gains from working, resulting in a utility of \$201.12 regardless of work requirements. At the upper bound, the utility cost offsets the income gains plus the SNAP benefits, resulting in a utility gain of \$201.12 from eliminating work requirements.

[37]This calculation assumes that earnings impacts are concentrated among participants who are induced to exit by work requirements. If we instead assume that earnings impacts are concentrated among participants who remain on SNAP, then eliminating work requirements increases the numerator by a further \$16.04, bounding the numerator between \$74.47+\$16.04 = \$90.51 and \$116.71+\$16.04 = \$132.75 (see discussion of benefit phase-out above).

[38]If the earnings response is driven by incentive effects, then the MVPF is bounded below by $\$90.51/\$99.12 = 0.91$ and bounded above by $\$116.71/\$93.91 = 1.18$ (see the discussion of benefit phase-out above).

would then be bounded between \$74.47/(\$72.05 + \$19.60) = \$74.47/\$91.65 = 0.81 and \$116.71/\$91.65 = 1.27. However, because these bounds are calculated from the upper bound of the confidence interval of our most optimistic earnings response estimates, we view the resulting MVPF bounds as less reliable than the bounds of 0.90 to 1.40 from the 24-month point estimates or the 1 to 1.40 from the 18-month point estimates.

Recall that most existing estimates of the MVPF of various aspects of SNAP and cash transfer programs are near 1 (Hendren and Sprung-Keyser 2020). Under the assumption that the earnings response to work requirements, if any, has a non-marginal utility cost, the MVPF of eliminating work requirements compares favorably with other policies targeting the SNAP population.