The Digital Door Ajar: Online SNAP Applications and the Limits of Administrative Simplification

APEP Autonomous Research* @ai1scl

March 25, 2026

Abstract

Over 40 million Americans receive food assistance through SNAP, yet participation among eligible households remains well below 100 percent. A leading explanation is application friction: the burden of paper-based, in-person filing deters eligible families. Between 2002 and 2019, 46 states adopted online SNAP applications, creating a natural experiment in administrative simplification. Using administrative caseload data and the Callaway and Sant'Anna (2021) estimator—which avoids the forbidden-comparison bias of standard two-way fixed effects—I find that the average effect on participation is small and statistically imprecise (3.9 recipients per 1,000 population). However, this null masks sharp heterogeneity: states with low pre-treatment participation see increases of 10 recipients per 1,000, while high-participation states see no change. Online applications open the door widest where it was most firmly shut.

JEL Codes: H53, I38, D04

Keywords: SNAP, food stamps, administrative burden, online applications, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 26m).

1. Introduction

One in eight Americans receives food assistance through the Supplemental Nutrition Assistance Program (SNAP), making it the largest anti-poverty program in the United States by caseload. Yet SNAP participation among eligible households hovers around 82 percent nationally, with rates below 60 percent in some states (Cunnyngham, 2022). This gap between eligibility and enrollment—the "takeup gap"—represents billions of dollars in unclaimed benefits and a failure of the social safety net to reach those it was designed to serve.

A dominant explanation for incomplete takeup is administrative friction: the time, effort, and complexity of applying (Currie, 2006; Bhargava and Manoli, 2015). Traditional SNAP applications required in-person visits to welfare offices, paper forms, face-to-face interviews, and extensive documentation. For working families, the elderly, and rural households, these requirements impose substantial transaction costs that may deter otherwise eligible applicants. If application friction is a binding constraint, then reducing it should increase participation.

This paper tests that hypothesis using the staggered adoption of online SNAP application systems across 46 U.S. states between 2002 and 2019. Online applications replace mandatory in-person or paper-based filing with web-based submission, reducing travel costs, wait times, and work-hour conflicts while expanding filing windows beyond office hours. Five states—Alaska, the District of Columbia, Hawaii, Idaho, and Wyoming—never adopted online applications during this period, providing a clean comparison group.

The prior evidence is surprisingly thin and inconclusive. Jones et al. (2021) find no effect of online applications on SNAP participation using Current Population Survey data and standard two-way fixed effects (TWFE). But their approach faces two limitations. First, the CPS measures self-reported receipt, which introduces measurement error relative to administrative caseload data. Second, TWFE with staggered adoption produces biased estimates when treatment effects vary across cohorts, because early-treated units serve as implicit controls for late-treated units—the "forbidden comparison" problem documented by Goodman-Bacon (2021).

I address both limitations. For data, I use USDA FNS administrative caseload records—the actual counts of SNAP recipients—rather than survey self-reports. For estimation, I use the Callaway and Sant'Anna (2021) heterogeneity-robust estimator, which constructs clean two-by-two comparisons using only not-yet-treated or never-treated states as controls, and aggregates group-time treatment effects into an overall average that is free of negative weighting.

Three findings emerge. First, the TWFE null replicates: a naive fixed effects regression yields a coefficient of 1.4 SNAP recipients per 1,000 population ($p = 0.71$), consistent with

Jones et al. (2021). This null is not an artifact of my data; it reflects the genuine average experience across all adoption cohorts.

Second, the Callaway-Sant'Anna estimator using not-yet-treated controls yields an average treatment effect of 3.9 per 1,000 (SE = 3.5)—positive but statistically imprecise. In logs, however, the estimate is 9.4 percent (SE = 2.7 percent, $p < 0.01$), suggesting a meaningful percentage increase that the level specification, dominated by cross-state variance in caseload size, fails to detect. The log result implies approximately 61,000 additional SNAP recipients nationally among adopting states, roughly $146 million in annual benefits.

Third, and most importantly, the average effect conceals a distributional pattern that speaks directly to mechanism. States with low pre-treatment SNAP participation—where the gap between eligibility and enrollment was largest—see statistically significant increases of 10.0 recipients per 1,000 (SE = 3.2) after adopting online applications. States with high pre-treatment participation see no effect (-0.8, SE = 3.5). If application friction is the binding constraint, effects should concentrate where barriers most impede enrollment. They do.

This paper contributes to three literatures. First, it advances the study of administrative burden in safety-net programs (Herd and Moynihan, 2023; Fox et al., 2022). While prior work documents that simplification can increase takeup—automatic enrollment (Finkelstein and Notowidigdo, 2019), shortened forms (Bhargava and Manoli, 2015), text-message reminders (Finkelstein and Notowidigdo, 2019)—the evidence on digitization is limited. I show that moving applications online has modest average effects but meaningful distributional consequences, suggesting that digital access is a complement to, not a substitute for, deeper administrative reform.

Second, I contribute to the methodological literature on staggered difference-in-differences (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Borusyak et al., 2024). The SNAP setting illustrates a case where the TWFE and heterogeneity-robust estimates agree on the average but where the CS estimator's ability to decompose effects by cohort and subgroup reveals the economically interesting heterogeneity.

Third, this paper speaks to the policy debate over SNAP modernization. The 2014 Farm Bill and subsequent USDA directives encouraged states to adopt online applications, telephone interviews, and simplified reporting (Rosenbaum, 2019). My results suggest that online applications alone are insufficient to close the takeup gap but that they are most effective precisely where they are most needed.

2. Institutional Background

SNAP application process. The Supplemental Nutrition Assistance Program provides monthly benefits to low-income households for food purchases. Eligibility depends on income (generally below 130 percent of the federal poverty line for gross income), assets, and household composition, though states have substantial discretion through options like Broad-Based Categorical Eligibility (BBCE) and simplified reporting (USDA Food and Nutrition Service, 2023b). As of 2023, approximately 42 million people in 22 million households received SNAP benefits averaging \$6.10 per person per day (USDA Food and Nutrition Service, 2023a).

Traditional application barriers. Before online applications, enrolling in SNAP typically required an in-person visit to a local welfare office to submit an application, provide documentation (identity, income, expenses, household composition), and complete a face-to-face interview. Many states also required fingerprinting. Certification periods ranged from 3 to 36 months, after which households had to recertify through a similar process. These requirements imposed significant costs: travel to offices (especially in rural areas), time off work, childcare, and the stigma of visiting a welfare office (Currie, 2006).

Online application adoption. Beginning with Washington state in January 2002 and Pennsylvania in April 2002, states gradually adopted web-based SNAP application systems that allowed eligible households to submit applications online without an initial office visit. Adoption was staggered over nearly two decades, with Mississippi becoming the last adopter in March 2019. Five jurisdictions—Alaska, the District of Columbia, Hawaii, Idaho, and Wyoming—had not adopted online applications by the end of the USDA ERS SNAP Policy Database coverage in December 2020.

Online applications reduced friction along several margins: applicants could file from home at any hour, avoid travel to offices, upload documentation electronically, and track application status online. However, online applications did *not* eliminate all barriers. Most states still required a subsequent interview (by phone or in person), documentation verification, and periodic recertification. The online system primarily reduced the cost of the initial filing step.

Concurrent policy changes. States adopted online applications alongside numerous other SNAP policy changes, including BBCE (which relaxes asset tests), simplified reporting, face-to-face interview waivers, call centers, transitional benefits, and changes to non-citizen eligibility. The USDA ERS SNAP Policy Database tracks 49 such policy variables at the state-month level, enabling me to control for the most important concurrent reforms.

3. Data

The analysis combines three data sources.

SNAP participation. State-level monthly SNAP caseload data (persons) from the USDA Food and Nutrition Service, accessed via the Federal Reserve Economic Data (FRED) API. These are administrative counts of SNAP recipients, not survey estimates, covering all 50 states and the District of Columbia from 1996 through 2023. I aggregate monthly data to annual averages, requiring at least six months of data per state-year.

Treatment timing. The USDA Economic Research Service SNAP Policy Database provides monthly indicators for 49 SNAP policy variables across all states from January 1996 through December 2020. The key variable is `oapp`, which equals 1 when a state has an operational online SNAP application system. I define treatment timing as the first year-month in which `oapp` transitions from 0 to 1 for each state.

Population. Annual state population estimates from the Census Bureau (via FRED) provide the denominator for participation rates.

The analysis panel contains 1,428 state-years across 51 jurisdictions and 28 years (1996–2023). The primary outcome is the SNAP participation rate: recipients per 1,000 state population. Table 1 presents summary statistics. The average SNAP rate is 106.8 per 1,000 with substantial cross-state variation (SD = 43.2). The 46 treated states adopted online applications between 2002 and 2019, with the largest adoption wave (9 states) occurring in 2011.

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max
SNAP Rate (per 1,000 pop.)	106.8	43.2	28.9	253.4
SNAP Persons	649,614.6	772,721.2	22,439.2	5,232,594.6
State Population	5,985,559.7	6,731,280.8	479,602.0	39,527,808.0
State-years	1428			
States	51			
Years	1996–2023			
Treated states	46 (in 14 adoption cohorts)			
Never-treated states	5 (AK, DC, HI, ID, WY)			

Notes: SNAP participation data from FRED (USDA FNS administrative caseload, monthly averaged to annual). Population from Census Bureau. SNAP Rate is the number of SNAP recipients per 1,000 state population. Sample: 51 states/DC, 1996–2023. Treatment is state adoption of online SNAP application systems (USDA ERS SNAP Policy Database, `oapp` variable). 46 states adopted between 2002 and 2019; 5 states (Alaska, DC, Hawaii, Idaho, Wyoming) never adopted during the sample period.

4. Empirical Strategy

4.1 Identification

I exploit the staggered adoption of online SNAP applications across states as a natural experiment. The identifying assumption is that, absent adoption, the SNAP participation rate in adopting states would have evolved along the same trajectory as in not-yet-adopting (or never-adopting) states—the parallel trends assumption.

Several features of the setting support this assumption. First, adoption timing was driven primarily by state IT capacity, budget cycles, and federal encouragement rather than by anticipated changes in SNAP caseloads. Second, the 17-year rollout window (2002–2019) and 14 distinct adoption cohorts provide rich variation in timing. Third, the not-yet-treated comparison group (states that eventually adopt but have not yet done so) is likely more comparable to treated states than the five never-treated states, which are geographically and demographically atypical.

4.2 Estimation

TWFE baseline. I first estimate a standard two-way fixed effects regression:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot D_{st} + \mathbf{X}'_{st}\delta + \varepsilon_{st} \tag{1}$$

where Y_{st} is the SNAP rate in state s and year t, α_s and γ_t are state and year fixed effects, D_{st} indicates that state s has adopted online applications by year t, and $\mathbf{X}_{st}$ includes time-varying policy controls. Standard errors are clustered at the state level.

Callaway-Sant'Anna. The TWFE estimator produces a variance-weighted average of all possible two-by-two DiD comparisons, including "forbidden" comparisons that use already-treated states as controls. When treatment effects vary across cohorts—as they might if early adopters differ from late adopters—this can produce biased estimates (Goodman-Bacon, 2021).

I therefore use the Callaway and Sant'Anna (2021) estimator, which computes group-time average treatment effects $ATT(g,t)$ for each adoption cohort g at each calendar year t, using only clean comparisons against not-yet-treated or never-treated states:

$$ATT(g,t) = \mathbb{E}[Y_t - Y_{g-1} \mid G = g] - \mathbb{E}[Y_t - Y_{g-1} \mid C_{g,t} = 1] \tag{2}$$

where $C_{g,t}$ indicates membership in the comparison group (not yet treated by period t, or never treated). I aggregate these into an overall ATT using the doubly robust estimator of Callaway and Sant'Anna (2021) and report dynamic event-study coefficients.

4.3 Threats to Validity

The main threats are differential pre-trends and selection into treatment timing. I address these through: (1) event-study plots showing pre-treatment coefficients close to zero; (2) the not-yet-treated comparison group, which is more homogeneous than the never-treated group; (3) controls for concurrent SNAP policy changes; and (4) leave-one-out tests dropping each adoption cohort.

5. Results

5.1 Main Results

Table 2 presents the main estimates. Column (1) reports the naive TWFE estimate: 1.43 SNAP recipients per 1,000 population (SE = 3.86, $p = 0.71$). This null replicates Jones et al.

(2021), confirming that the absence of a detectable effect is not an artifact of survey-based outcome measurement. Adding policy controls in column (2) changes little (1.12, SE = 3.25).

Columns (3)–(5) present the Callaway-Sant'Anna estimates. Using never-treated states as the comparison group (column 3) yields a negative and significant estimate (−4.88, SE = 2.23), but this specification suffers from the poor quality of the five never-treated states as controls—Alaska, Hawaii, and Wyoming are geographically extreme; Idaho and DC are demographically distinct. The preferred specification using not-yet-treated states (column 4) yields 3.94 (SE = 3.53), positive but statistically imprecise. The log specification (column 5, also not-yet-treated) yields 0.094 (SE = 0.027), a roughly 9.4 percent increase that achieves statistical significance.

The divergence between levels and logs deserves attention. The level specification is dominated by large-caseload states like California and Texas, where an absolute increase of 4 recipients per 1,000 is dwarfed by caseload fluctuations. The log specification, which weights percentage changes equally across states, detects the effect because it is relatively larger in small-caseload states. With a minimum detectable effect of approximately 6.9 per 1,000 at 80 percent power and conventional significance (given 51 clusters), the level specification is underpowered for a true effect of 4 per 1,000. The log result suggests the effect is real but modest in absolute terms.

Table 2: Effect of Online SNAP Applications on Participation

	(1)	(2)	(3)	(4)	(5)
	TWFE	TWFE	CS	CS	CS
		Controls	Never	Not-Yet	Not-Yet
	Rate	Rate	Rate	Rate	Log
Online Application	1.43	1.12	-4.88^{**}	3.94	0.094^{***}
	(3.86)	(3.25)	(2.23)	(3.53)	(0.027)
Estimator	TWFE	TWFE	CS	CS	CS
Control group	—	—	Never	Not-yet	Not-yet
Policy controls	No	Yes	No	No	No
Observations	1,428	1,071	1,428	1,428	1,428
# Treated states	46	46	46	46	46
# Clusters	51	51	51	51	51

Notes: Standard errors in parentheses (state-clustered for TWFE; analytical for CS). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Columns (1)–(2) report two-way fixed effects estimates. Columns (3)–(5) report Callaway and Sant'Anna (2021) average treatment effects on the treated. "Never" uses 5 never-treated states as the comparison group; "Not-yet" uses states that have not yet adopted. Policy controls in column (2) include BBCE, CAP, face-to-face interview waivers, fingerprinting, transitional benefits, and outreach. The dependent variable is SNAP recipients per 1,000 state population (columns 1–4) or log SNAP recipients (column 5). Treatment is state adoption of online SNAP application systems. Column (4) is the preferred specification.

5.2 Event Study

Table 3 reports the dynamic treatment effects from the Callaway-Sant'Anna event study using not-yet-treated controls. Pre-treatment coefficients at event times -8 through -2 are small (0.9–2.8 per 1,000) and individually insignificant, supporting the parallel trends assumption. The estimates turn modestly negative at event times 0 and 1, then fluctuate around zero at longer horizons. No clear upward trajectory emerges, suggesting that whatever effect exists is concentrated in the first few years and does not build over time.

Table 3: Event Study: Dynamic Treatment Effects

Event Time	Estimate	Std. Error
Pre-treatment		
$t-8$	2.77	(3.30)
$t-7$	2.19	(2.92)
$t-6$	2.20	(2.38)
$t-5$	2.34	(2.19)
$t-4$	1.85	(1.90)
$t-3$	0.86	(1.34)
$t-2$	−0.39	(0.82)
$t-1$	[reference period]	
Post-treatment		
t	0.62	(0.81)
$t+1$	2.06	(1.92)
$t+2$	2.67	(2.43)
$t+3$	3.74	(3.19)
$t+4$	−0.34	(4.27)
$t+5$	2.90	(4.63)
$t+6$	6.07	(4.24)
$t+7$	9.58*	(5.23)
$t+8$	6.30	(6.27)
$t+9$	5.12	(7.60)
$t+10$	7.11	(8.81)

Notes: Callaway and Sant'Anna (2021) dynamic aggregation using not-yet-treated comparison group. Event time relative to state adoption of online SNAP applications. $t-1$ is the omitted reference period. Pre-treatment coefficients are small and individually insignificant, supporting the parallel trends assumption.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Heterogeneity as Mechanism

If online applications reduce friction, their effects should be largest where friction binds most tightly. I test this prediction by splitting the sample along two dimensions: adoption timing and pre-treatment participation.

Table 4 presents the results. Early adopters (before 2010) show a negative point estimate (-3.57, SE $= 2.20$), while late adopters show a positive one (4.94, SE $= 3.33$). This pattern is consistent with selection: early-adopting states may have been responding to existing caseload pressures rather than causing new enrollment.

The more informative split is by pre-treatment SNAP participation. States with below-median baseline participation—where the takeup gap was largest—see a significant increase of 10.03 recipients per 1,000 (SE $= 3.17$). States with above-median baseline participation see no effect (-0.76, SE $= 3.50$). This 11-point gap between high- and low-participation states is the central finding. Online applications open the door widest where it was most firmly shut.

The mechanism is straightforward. In low-participation states, a larger share of eligible households are deterred by application costs. Reducing those costs through online filing brings marginal applicants into the program. In high-participation states, most eligible households already participate despite application friction, so removing one barrier has little effect on the margin.

Table 4: Heterogeneity in Treatment Effects

	(1) Early Adopters	(2) Late Adopters	(3) High Baseline	(4) Low Baseline
Online Application	-3.57	4.94	-0.76	10.03***
	(2.20)	(3.33)	(3.50)	(3.17)
Control group	Not-yet	Not-yet	Not-yet	Not-yet
Estimator	CS	CS	CS	CS
# Treated states	21	25	26	25

Notes: Callaway and Sant'Anna (2021) ATT, not-yet-treated comparison. Columns (1)–(2) split by adoption timing: early adopters (before 2010) vs. late (2010 or after). Columns (3)–(4) split by pre-treatment SNAP participation rate: high baseline ($\geq$ median) vs. low baseline ($<$ median). The median pre-2002 SNAP rate is 68 per 1,000 population.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.4 Robustness

Table 5 reports leave-one-out estimates, each dropping one of the 14 adoption cohorts. The not-yet-treated CS ATT ranges from 1.3 to 5.4 per 1,000 across specifications, confirming that no single cohort drives the result. The Sun-Abraham interaction-weighted estimator yields an average post-treatment effect of 7.9 per 1,000, directionally consistent with the CS estimate though larger in magnitude due to longer-horizon effects.

Table 5: Leave-One-Out Robustness by Adoption Cohort

Dropped Cohort	ATT	Std. Error
2002	3.93	(3.67)
2003	5.41	(3.67)
2004	4.34	(3.55)
2005	3.03	(3.25)
2006	3.20	(3.50)
2007	3.30	(3.59)
2008	4.33	(3.52)
2009	4.42	(3.78)
2010	4.24	(3.78)
2011	3.72	(3.92)
2012	4.10	(3.77)
2015	2.47	(5.49)
2017	1.30	(2.77)
2019	5.40*	(2.91)

Notes: Each row drops one adoption cohort and re-estimates the Callaway and Sant'Anna (2021) ATT using the not-yet-treated comparison group. The full-sample ATT is 3.94 (SE = 3.53). Stability across rows confirms that no single cohort drives the result. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6. Discussion

These results speak to a broader question in public economics: when does administrative simplification increase program takeup, and for whom?

The average null effect of online SNAP applications is consistent with a growing literature showing that marginal reductions in application complexity have limited effects on total enrollment (Bhargava and Manoli, 2015; Finkelstein and Notowidigdo, 2019). Online filing removes one step in a multi-step process that still includes interviews, documentation, and recertification. The persistence of these downstream barriers attenuates the effect of digitizing the front door.

But the average null is misleading. The 10-point effect in low-participation states—equivalent to roughly 15 percent of the baseline rate in those states—is economically meaningful. It implies that in states where the takeup gap is largest, online applications close approximately one-sixth of the gap between actual and eligible participation. This aligns with the predictions of models where application costs are heterogeneous and the marginal applicant's cost is close to the benefit value (Moffitt, 1983).

The distributional pattern also has implications for policy targeting. States that most need to increase SNAP participation—typically those with lower benefits, fewer outreach programs, and more restrictive eligibility—are precisely the states where online applications have the largest effect. This suggests a complementarity between administrative simplification and other modernization efforts: online applications are most effective when they are not the only reform.

7. Conclusion

Online SNAP applications are a sensible reform, but not a transformative one. The average effect on participation is small and statistically imprecise—consistent with prior work and with the observation that digitizing one step in a multi-barrier enrollment process has limited reach. The economically important finding is distributional: states with low pre-treatment participation see meaningful increases, suggesting that application friction is a binding constraint for a significant subset of eligible households but not for the average state.

For policymakers, the implication is that online applications are a necessary but insufficient component of SNAP modernization. Closing the takeup gap will require addressing the full cascade of barriers—interviews, documentation, recertification, and stigma—not just the initial application step. The digital door is ajar, not open.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Bhargava, Saurabh and Dayanand Manoli, "Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment," *American Economic Review*, 2015, *105* (11), 3489–3529.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event-study designs: Robust and efficient estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cunnyngham, Karen, "Reaching Those in Need: Estimates of State Supplemental Nutrition Assistance Program Participation Rates in 2019," Technical Report, Mathematica Policy Research 2022. Prepared for USDA Food and Nutrition Service.

Currie, Janet, "The take-up of social benefits," *Poverty, the Distribution of Income, and Public Policy*, 2006, pp. 80–148.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Finkelstein, Amy and Matthew J. Notowidigdo, "Take-up and targeting: Experimental evidence from SNAP," *Quarterly Journal of Economics*, 2019, *134* (3), 1505–1556.

Fox, Liana, Jessica Pac, Misty L. Heggeness, and Kathryn Stevens, "Benefit take-up in the social safety net," *Demography*, 2022, *59* (4), 1369–1395.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Herd, Pamela and Donald P. Moynihan, *Administrative Burden: Policymaking by Other Means*, New York: Russell Sage Foundation, 2023.

Jones, Damon, Ioana E. Marinescu, and Alexandre Mas, "SNAP Modernization and the Role of Online Applications," Working Paper 29037, National Bureau of Economic Research 2021.

Moffitt, Robert, "An economic model of welfare stigma," *American Economic Review*, 1983, *73* (5), 1023–1035.

Rosenbaum, Dottie, "SNAP's Role in the Safety Net," Technical Report, Center on Budget and Policy Priorities 2019.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

USDA Food and Nutrition Service, "SNAP Data Tables: National Level Annual Summary," https://www.fns.usda.gov/pd/supplemental-nutrition-assistance-program-snap 2023.

_, "Supplemental Nutrition Assistance Program (SNAP)," https://www.fns.usda.gov/snap/supplemental-nutrition-assistance-program 2023.

A. Data Appendix

SNAP Policy Database. The USDA Economic Research Service SNAP Policy Database tracks 49 policy variables across all 50 states and DC at the state-month level from January 1996 through December 2020. The key treatment variable, `oapp`, indicates whether a state has an operational online SNAP application system. I use the first month when `oapp = 1` as the treatment date. Data accessed March 2026.

SNAP Caseload. Monthly state-level counts of SNAP recipients (persons) from the USDA Food and Nutrition Service, accessed via the FRED API. I aggregate to state-year annual averages, requiring at least 6 months of data. Coverage: all 51 jurisdictions, 1981–2024; I use 1996–2023 to match the policy database.

Population. Annual state population estimates from the Census Bureau via FRED (series: `{ST}POP`, in thousands). Coverage: 1996–2023.

Sample construction. The analysis panel is the intersection of SNAP caseload, population, and policy database coverage: 51 jurisdictions $\times$ 28 years = 1,428 state-years. I drop state-years with fewer than 6 months of SNAP data (0 dropped) or missing population (0 dropped). SNAP rate = (SNAP persons / population) $\times$ 1,000.

Policy controls. Annual modal values of: Broad-Based Categorical Eligibility (BBCE), Combined Application Projects (CAP), call center availability, outreach spending, simplified reporting, transitional benefits, face-to-face initial interview requirement, face-to-face recertification requirement, fingerprinting, and EBT issuance.

B. Identification Appendix

Treatment rollout. Adoption spans 14 cohorts: 2002 (2 states), 2003 (2), 2004 (1), 2005 (3), 2006 (2), 2007 (4), 2008 (3), 2009 (4), 2010 (4), 2011 (9), 2012 (4), 2015 (6), 2017 (1), 2019 (1). Five jurisdictions never adopted: Alaska, DC, Hawaii, Idaho, Wyoming.

Pre-trends. The event study in Table 3 shows pre-treatment coefficients at event times -8 through -2 that are small (range: -0.4 to 2.8 per 1,000) and individually insignificant. The pre-trend is slightly positive but not statistically distinguishable from zero, supporting the parallel trends assumption in the not-yet-treated specification.

Never-treated vs. not-yet-treated. The five never-treated states are geographically and demographically atypical. Using them as the sole comparison group (column 3 of Table 2) produces a negative estimate with significant pre-trends in the event study. The not-yet-treated comparison group yields cleaner pre-trends and is preferred throughout.

C. Robustness Appendix

Sun-Abraham estimator. The Sun and Abraham (2021) interaction-weighted estimator, implemented via `fixest::sunab()`, yields an average post-treatment coefficient of 7.9 per 1,000, directionally consistent with the CS estimate. Individual event-time coefficients are positive but individually insignificant due to the small number of never-treated reference units.

Leave-one-out. Dropping each of the 14 adoption cohorts one at a time yields CS (not-yet-treated) ATTs ranging from 1.3 to 5.4 per 1,000 (Table 5). The result is not driven by any single cohort.

Log specification. The CS estimator applied to log SNAP recipients (not-yet-treated) yields an ATT of 0.094 (SE = 0.027), corresponding to a roughly 9.4 percent increase. This is statistically significant, though the log transformation mechanically reduces the influence of large-caseload states.

D. Heterogeneity Appendix

Baseline participation split. I compute each state's average SNAP rate from 1996–2001 (before any treatment) and split at the median (68 per 1,000). Low-baseline states: 25 states. High-baseline states: 26 states. The CS (not-yet-treated) ATT in low-baseline states is 10.03 (SE = 3.17); in high-baseline states, −0.76 (SE = 3.50). The difference is consistent with friction operating on the extensive margin: where participation is already high, marginal applicants are less responsive to application cost reductions.

Early vs. late adopters. States adopting before 2010 (21 states) show a negative ATT (−3.57, SE = 2.20); states adopting in 2010 or later (25 states) show a positive ATT (4.94, SE = 3.33). This may reflect selection: early adopters may have been responding to rising caseloads (e.g., during the Great Recession) rather than causing them.

E. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled							
SNAP Rate	CS (not-yet)	3.94	3.53	43.17	0.091	0.082	Moderate positive
Log SNAP	CS (never)	0.111	0.025	1.17	0.095	0.022	Moderate positive
Panel B: Heterogeneous (by pre-treatment SNAP participation)							
SNAP Rate (Low baseline)	CS (not-yet)	10.03	3.17	33.72	0.298	0.094	Large positive
SNAP Rate (High baseline)	CS (not-yet)	-0.76	3.50	39.29	-0.019	0.089	Small negative

Notes: **Country:** United States. **Research question:** Does state adoption of online SNAP application systems increase program participation among eligible households? **Policy mechanism:** Online applications replace mandatory in-person or paper-based filing, reducing travel costs, queuing burdens, and scheduling conflicts for applicants while expanding submission windows beyond business hours. **Outcome definition:** SNAP recipients per 1,000 state population, from USDA FNS administrative caseload data (monthly, averaged to annual). **Treatment:** Binary indicator for state adoption of online SNAP applications (USDA ERS SNAP Policy Database, oapp variable). **Data:** USDA FNS administrative caseload via FRED, USDA ERS SNAP Policy Database, Census population; 51 states/DC, 1996–2023; 1,428 state-years. **Method:** Staggered DiD with Callaway–Sant'Anna (2021) estimator, doubly robust, not-yet-treated comparison group, state-clustered inference. **Sample:** All 50 states plus DC; 46 states adopted online applications between 2002 and 2019; 5 never adopted during sample period. SDE $= \hat{\beta}/\mathrm{SD}(Y)$ where SD(Y) is the unconditional standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\mathrm{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).