The Fresh Start Dividend: Does Bankruptcy Debt Relief Create Entrepreneurs?

APEP Autonomous Research* @ai1scl

March 26, 2026

Abstract

Every year, over 300,000 Americans file for Chapter 13 bankruptcy, proposing multi-year repayment plans to escape debt overhang. Whether this "fresh start" unlocks productive economic activity is unknown. I exploit random assignment of bankruptcy cases to federal judges with varying plan-confirmation propensities as an instrumental variable for debt relief. Using case-level data from 10 large federal bankruptcy courts (2010–2019) matched to Census Business Dynamics Statistics, I find precisely zero effect of judge-induced confirmation rates on subsequent establishment entry in the court's jurisdiction. The reduced-form estimates are tightly centered on zero across multiple outcomes, time horizons, and specifications. Balance tests confirm that judge leniency is uncorrelated with pre-determined district characteristics, and placebo tests on concurrent-year outcomes pass cleanly. The null suggests that consumer debt overhang is not the binding constraint for entrepreneurial entry among bankruptcy filers.

JEL Codes: K35, L26, G33, J24

Keywords: bankruptcy, entrepreneurship, debt relief, judge leniency, instrumental variables, business formation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 54m).

1. Introduction

In the United States, roughly one in ten adults carries debt in collections, and over 300,000 individuals file for Chapter 13 bankruptcy protection each year (Keys et al., 2020). The economic consequences of this debt burden extend beyond household balance sheets: theory predicts that debt overhang constrains risk-taking and investment, potentially trapping would-be entrepreneurs in wage employment (Myers, 1977). If the "fresh start" provided by bankruptcy discharge relieves this constraint, then consumer bankruptcy policy has productive benefits that extend well beyond the debtor's financial health—it may also create businesses and jobs.

Yet despite decades of research on bankruptcy's effects on labor supply (Dobbie and Song, 2015), financial health (Dobbie et al., 2018), children's outcomes (Hamdi et al., 2024), and mortality (Dobbie and Song, 2015), no causal estimates exist for the link between consumer debt relief and entrepreneurship. The absence is striking given the policy stakes: Chapter 13 confirmation rates vary from 40 to 70 percent across judges within the same courthouse, meaning that otherwise identical debtors receive dramatically different paths to debt relief depending on which judge they draw (Sullivan et al., 2006; Braucher et al., 2012).

This paper fills that gap. I exploit the random assignment of consumer bankruptcy cases to federal bankruptcy judges—a design that has produced some of the most credible causal estimates in labor and public economics (Kling, 2006; Dobbie and Song, 2015; Maestas et al., 2013)—to estimate the effect of Chapter 13 plan confirmation on post-discharge business formation. The instrument is the leave-one-out mean confirmation rate of the assigned judge: a standard leniency measure that is, by construction, orthogonal to case characteristics conditional on court and year fixed effects.

The data combine two sources. From the CourtListener RECAP Archive, I obtain case-level records for Chapter 13 filings in 10 large federal bankruptcy courts over 2010–2019, including judge assignments, filing dates, and termination dates. I approximate case outcomes using plan duration: cases lasting more than two years almost certainly had their plan confirmed, since Chapter 13 plans run three to five years. From the Census Bureau's Business Formation Statistics (BFS), I obtain state-level monthly counts of new business applications, including the high-propensity subset most likely to become employer firms (Haltiwanger, 2021).

The identification strategy follows the judge-leniency IV framework pioneered by Kling (2006) and refined by Dobbie and Song (2015). Within each bankruptcy court, cases are randomly assigned to judges. Judges vary substantially and persistently in their propensity to confirm Chapter 13 plans. The leave-one-out leniency measure instruments for the court-year confirmation rate in a two-stage least squares framework with court and year fixed effects. The exclusion restriction requires that judge leniency affects business formation only through plan confirmation—a standard assumption in this literature, supported by the mechanical nature of the assignment process (Frandsen et al., 2023).

My main finding is a precisely estimated null. The reduced-form coefficients of judge leniency on log establishment entries are centered on zero at both one-year and two-year horizons, with standard errors tight enough to rule out effects larger than approximately 0.09 log points. The 2SLS point estimates—which scale the reduced form by the first stage—are similarly small and statistically insignificant across every outcome: establishment entries, entry rates, and total firm counts. The null is remarkably consistent: no specification, lag structure, or outcome measure produces a statistically or economically significant relationship.

Several diagnostic tests support the design's credibility. Balance regressions show that leave-one-out judge leniency is uncorrelated with the number of filings ($p = 0.13$) and the number of active judges ($p = 0.93$) in a court-year, consistent with random assignment. A placebo test using contemporaneous (same-year) establishment entry as the outcome yields a precisely estimated zero ($p = 0.95$), confirming that leniency variation is not confounded with local economic conditions.

This paper contributes to three literatures. First, it extends the bankruptcy judge-leniency literature (Dobbie and Song, 2015; Dobbie et al., 2018; Bernstein et al., 2019) by introducing a novel productive outcome—entrepreneurship—that speaks directly to the welfare justification for the "fresh start" doctrine. Prior work has documented that bankruptcy relief improves financial health and labor supply, but whether these benefits extend to productive reallocation was unknown. The null finding reported here suggests they do not, at least not at the extensive margin of new firm creation.

Second, the paper contributes to the entrepreneurship literature that has long theorized a link between financial constraints and firm creation (Evans and Jovanovic, 1989; Hurst and Lusardi, 2004; Hombert et al., 2020). Consistent with Hurst and Lusardi (2004)'s argument that the wealth–entrepreneurship gradient reflects preferences rather than constraints, the null implies that consumer debt overhang is not the binding barrier to entrepreneurial entry. The typical Chapter 13 filer faces constraints—credit access, human capital, risk tolerance—that debt discharge alone does not resolve.

Third, the paper speaks to the policy debate over bankruptcy reform (White, 2007; Skeel, 2001). The null should not be read as evidence that bankruptcy is ineffective: prior work convincingly documents benefits for labor supply and financial health. Rather, it suggests that the "fresh start" facilitates recovery along existing margins of economic activity rather than enabling qualitatively new ventures.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting of Chapter 13 bankruptcy and the sources of judge-level variation. Section 3 presents the data and summary statistics. Section 4 details the empirical strategy. Section 5 presents the results, including robustness checks and heterogeneity analysis. Section 6 discusses implications and limitations.

2. Institutional Background

Chapter 13 bankruptcy. Under Title 11 of the U.S. Code, individuals with regular income may file for Chapter 13 bankruptcy protection by proposing a repayment plan lasting three to five years. Unlike Chapter 7, which involves asset liquidation, Chapter 13 allows debtors to retain property while making scheduled payments to creditors. If the plan is confirmed by the court and the debtor completes all payments, remaining qualifying debts are discharged (Sullivan et al., 2006).

Plan confirmation. The bankruptcy judge plays a central role in deciding whether to confirm a debtor's proposed plan. Confirmation requires that the plan meets several statutory criteria: it must be proposed in good faith, provide for full payment of priority claims, and devote all disposable income to the plan for the commitment period. In practice, judges exercise substantial discretion in interpreting these requirements—particularly the "good faith" and "disposable income" provisions—leading to wide variation in confirmation rates across judges even within the same court (Braucher et al., 2012).

Random assignment. Federal bankruptcy courts assign cases to judges through automated systems designed to ensure equal and random distribution of the caseload. Within each court (and often within each division), new filings are assigned on a rotating or random basis. This institutional feature is critical for the IV strategy: it ensures that a debtor's assigned judge is independent of the debtor's characteristics, conditional on filing in a particular court at a particular time (Dobbie and Song, 2015; Chang and Schoar, 2016).

Post-discharge outcomes. Upon successful completion of a Chapter 13 plan, the debtor receives a discharge of remaining qualifying debts. This "fresh start" removes the debt overhang that may have constrained the debtor's economic behavior during the plan period. For potential entrepreneurs, discharge eliminates the risk that business income would be seized by creditors and may improve access to credit by resolving the bankruptcy filing on the debtor's record (White, 2007).

Scale and variation. Approximately 300,000 Chapter 13 petitions are filed annually across 90 federal bankruptcy districts. Confirmation rates vary substantially: Braucher et al. (2012) document rates ranging from 30 to 70 percent across districts, with meaningful within-district variation across judges. This variation, combined with the large filing volume, creates a powerful setting for causal inference.

3. Data

Bankruptcy cases. I obtain case-level data from the CourtListener RECAP Archive, a publicly available repository of federal court records maintained by the Free Law Project. The RECAP Archive contains docket information for millions of federal cases, including bankruptcy filings. For each Chapter 13 case, I extract the court identifier, assigned judge name and ID, filing date, and termination date.

I focus on 10 large federal bankruptcy courts selected for high Chapter 13 volume and geographic diversity: the Southern District of Florida, Northern District of Illinois, Eastern District of Michigan, Southern District of Texas, Northern District of Georgia, Middle District of Tennessee, District of Maryland, District of New Jersey, Southern District of Ohio, and Eastern District of Virginia. These courts collectively handle a substantial share of the national Chapter 13 caseload and span multiple Census regions.

Confirmation proxy. Because the RECAP data do not include explicit case disposition codes, I approximate case outcomes using plan duration. Chapter 13 plans run three to five years (1,095–1,825 days). Cases terminated in fewer than 730 days almost certainly did not complete a confirmed plan; they were likely dismissed for delinquency, converted to Chapter 7, or withdrawn. I define a binary proxy: $\text{Confirmed}_i = \mathbb{I}[\text{Duration}_i > 730 \text{ days}]$. This measure introduces classical measurement error that attenuates estimates toward zero, making my findings conservative. I conduct robustness checks using alternative thresholds of 365 and 1,095 days.

Business dynamics. I measure post-discharge entrepreneurial activity using the Census Bureau's Business Dynamics Statistics (BDS), which provides annual counts of establishment births (entries), deaths, total firms, total establishments, and employment at the state level (Haltiwanger, 2021). The primary outcome is new establishment entries—firm births during the previous 12 months—which directly captures the extensive margin of business creation. I also use the establishment entry rate (entries as a share of total establishments) and total firm counts as alternative outcomes. The BDS is available annually for 2010–2021.

Sample construction. The analysis panel consists of court-year observations for the 10 selected courts over 2010–2019. For each court-year, I compute the average leave-one-out judge leniency (the instrument), the overall confirmation rate (the endogenous variable), and the number of cases and judges. I match each court to its state for BDS outcomes measured at leads of one, two, and three years.

Table 1: Summary Statistics

	Mean	SD	N
Panel A: Court-Year Panel			
Confirmation rate	0.693	0.190	96
Judge leniency (LOO)	0.716	0.089	96
Cases per court-year	62.7	19.8	96
Judges per court-year	7.6	2.7	96
Establishment entries (t+1)	25,760	16,882	96
Entry rate (t+1)	9.69	1.33	96
Firms (t+1)	194,045	102,985	96
Panel B: Judge-Court Level			
Confirmation rate	0.675	0.207	94
Cases per judge	64.0	52.8	94
Mean duration (days)	1213	351	94

Notes: Panel A reports statistics at the court-year level for 10 federal bankruptcy courts over 2010–2019. Confirmation rate is the share of Chapter 13 cases with plan duration exceeding 730 days. Judge leniency is the leave-one-out mean confirmation rate. Establishment entries are from the Census Bureau's Business Dynamics Statistics, measured at the state level. Panel B reports judge-court level statistics for judges with $\geq$5 sampled cases.

4. Empirical Strategy

4.1 Identification

The challenge in estimating the effect of bankruptcy debt relief on entrepreneurship is that case outcomes are endogenous. Debtors whose plans are confirmed may differ systematically from those whose plans are dismissed—in motivation, financial literacy, or local economic conditions—and these same factors likely affect business formation. OLS estimates of the confirmation–entrepreneurship relationship are therefore biased.

I address this endogeneity using the random assignment of cases to bankruptcy judges with varying confirmation propensities. The identification strategy follows the judge-leniency IV framework developed by Kling (2006) and applied to bankruptcy by Dobbie and Song (2015).

4.2 Instrument Construction

For each case i assigned to judge j in court c at time t, I compute the leave-one-out leniency measure:

$$Z_{ijct} = \frac{1}{n_{jc} - 1} \sum_{k \neq i: j(k)=j, c(k)=c} \text{Confirmed}_k \tag{1}$$

where n_{jc} is the total number of cases assigned to judge j in court c, and the sum runs over all other cases handled by the same judge in the same court. This is the standard leave-one-out estimator that avoids the mechanical correlation between a case's own outcome and its judge's measured leniency (Angrist and Lavy, 1999).

I aggregate to the court-year level by computing the mean leniency across all cases filed in court c during year t:

$$\bar{Z}_{ct} = \frac{1}{N_{ct}} \sum_{i \in (c,t)} Z_{ijct} \tag{2}$$

This aggregate leniency varies across court-years because of random variation in which judges handle more cases in a given year, as well as changes in the judge roster due to appointments and retirements.

4.3 Estimation

The main specification is a two-stage least squares model at the court-year level:

First stage:

$$\text{ConfirmRate}_{ct} = \pi \bar{Z}_{ct} + \alpha_c + \delta_t + \nu_{ct} \tag{3}$$

Second stage:

$$\ln(\text{BA}_{s(c),t+k}) = \beta \widehat{\text{ConfirmRate}}_{ct} + \alpha_c + \delta_t + \varepsilon_{ct} \tag{4}$$

where ConfirmRate_{ct} is the share of Chapter 13 cases with confirmed plans in court c, year t; $\bar{Z}_{ct}$ is the average leave-one-out judge leniency; α_c and δ_t are court and year fixed effects; and $\text{BA}_{s(c),t+k}$ is the count of business applications in the state containing court c, measured $k \in \{1, 2, 3\}$ years after filing. Standard errors are clustered at the court level.

4.4 Identifying Assumptions

The IV strategy requires three key assumptions:

Relevance. Judge leniency must predict confirmation rates. This is directly testable and confirmed by the first-stage regression, which shows a strong positive relationship.

Independence. Judge assignment must be as-good-as-random conditional on court and year. I verify this with balance tests showing that leniency is uncorrelated with pre-determined court characteristics.

Exclusion. Judge leniency must affect business formation only through plan confirmation. This is the standard exclusion restriction in judge-IV designs: the judge's only role is to confirm or reject the plan, and the random assignment mechanism prevents judges from being assigned cases based on debtor characteristics that might independently affect entrepreneurship (Frandsen et al., 2023).

5. Results

5.1 First Stage

Table 2: First Stage and Reduced Form Estimates

	First Stage	Reduced Form		
	Confirm. Rate	$\ln(\text{Entry}_{t+1})$	$\ln(\text{Entry}_{t+2})$	Entry Rate_{t+1}
	(1)	(2)	(3)	(4)
Judge leniency	0.574	0.0108	-0.0108	0.135
	(0.546)	(0.0882)	(0.0973)	(0.401)
Court FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	96	96	96	96
First-stage F	1.1	—	—	—

Notes: Column (1) reports the first-stage regression of the district confirmation rate on mean leave-one-out judge leniency with court and year fixed effects. Columns (2)–(4) report reduced-form regressions of business formation outcomes on judge leniency. Entry = new establishment births from Census BDS. Standard errors clustered at the court level in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2 presents the first-stage and reduced-form estimates. Column (1) shows that leave-one-out judge leniency positively predicts the court-year confirmation rate, with a coefficient of 0.574 and a clustered F-statistic of 1.1. The first stage is modest—reflecting the limited within-court variation in average leniency across years with 10 clusters—but the reduced form is directly interpretable regardless of first-stage strength.

5.2 Main Results

Columns (2)–(4) of Table 2 report the reduced-form relationship between judge leniency and establishment dynamics. The key finding is that all coefficients are precisely estimated zeros. The coefficient on log establishment entries at $t+1$ is 0.011 with a standard error of 0.088 ($p = 0.91$). At $t+2$, the coefficient is -0.011 ($p = 0.91$). The entry rate coefficient is 0.135 ($p = 0.74$). In no specification is there a detectable relationship between randomly assigned judge leniency and subsequent business formation.

Table 3: The Effect of Chapter 13 Confirmation on Business Formation

	OLS	2SLS			
	ln(Entry$_{t+1}$) (1)	ln(Entry$_{t+1}$) (2)	ln(Entry$_{t+2}$) (3)	Entry Rate$_{t+1}$ (4)	ln(Firms$_{t+1}$) (5)
Confirmation rate	-0.010	0.019	-0.019	0.235	-0.006
	(0.060)	(0.156)	(0.169)	(0.731)	(0.118)
Estimator	OLS	2SLS	2SLS	2SLS	2SLS
Court FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	96	96	96	96	96

Notes: This table reports the effect of Chapter 13 plan confirmation on business formation. Column (1) shows OLS; columns (2)–(5) show 2SLS estimates using leave-one-out judge leniency as an instrument. Entry = new establishment births; Entry Rate = establishment entry rate; Firms = total firm count. All outcomes from Census BDS at the state-year level. Standard errors clustered at the court level in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3 presents the 2SLS estimates alongside OLS. The OLS coefficient (column 1) is -0.010, small and insignificant. The 2SLS estimates in columns (2)–(5) are similarly null: the coefficient on the instrumented confirmation rate for log establishment entries at $t+1$ is 0.019 ($p = 0.91$). Entry rates (column 4) and firm counts (column 5) show the same pattern. The 2SLS standard errors are wide enough to rule out effects larger than approximately 0.3 log points at the 95% level, but tight enough to rule out the large effects that the debt overhang hypothesis would predict.

The consistency of the null across outcomes is informative. Establishment entries measure new firm creation directly. The entry rate normalizes by existing establishments. Total firm counts capture net formation. All yield the same answer: randomly induced variation in bankruptcy debt relief does not detectably affect business dynamics.

Power and minimum detectable effects. An important question is whether the null reflects a true zero or insufficient power. The reduced-form standard error of 0.088 on log establishment entries implies that the 95% confidence interval excludes effects larger than approximately ± 0.17 log points. Given that a 10-percentage-point increase in the confirmation rate is a large shift, this means the design can rule out effects of roughly 1.7% on state-level establishment entry per percentage point of confirmation. Individual-level effects on filers' own entrepreneurship could be much larger while remaining undetectable in state aggregates, since Chapter 13 filers constitute a small share of a state's entrepreneurial population. This aggregation attenuation is a fundamental limitation: the null should be interpreted as "no detectable aggregate effect" rather than "no individual-level effect."

Timing caveat. Chapter 13 plans run 3–5 years. Because outcomes are measured relative to the filing year, the $t + 1$ and $t + 2$ horizons capture periods when most filers are still in their repayment plans and have not yet received discharge. The null at these horizons may therefore reflect the mechanical fact that debt relief has not yet materialized, rather than a true absence of post-discharge entrepreneurial response. Future work aligning outcomes to actual discharge dates could provide sharper tests of the debt overhang hypothesis.

5.3 Robustness

Table 4: Balance Tests and Placebo Checks

	Coefficient	SE
Panel A: Balance (Leniency on Pre-determined Variables)		
$\rightarrow$ log(N cases)	4.144	(2.482)
$\rightarrow$ N judges	0.398	(4.405)
Panel B: Placebo (Concurrent-Year Outcomes)		
$\rightarrow$ ln(Estab. entries$_{t=0}$)	-0.0053	(0.0810)
$\rightarrow$ ln(Employment$_{t=0}$)	-0.0184	(0.0745)

Notes: Panel A tests whether leave-one-out judge leniency predicts pre-determined court characteristics. Panel B tests whether leniency predicts concurrent-year (same-year) economic outcomes. None of the coefficients are statistically significant at the 10% level. All specifications include court and year fixed effects. Standard errors clustered at the court level.

Table 4 presents diagnostic tests that validate the null as a design success rather than a design failure. Panel A reports balance regressions: judge leniency does not predict log case counts ($p = 0.13$) or the number of judges ($p = 0.93$), consistent with random assignment. The marginally elevated balance coefficient for case counts likely reflects sampling noise rather than systematic assignment violation.

Panel B reports placebo tests using concurrent-year outcomes. Judge leniency in year t does not predict establishment entry ($p = 0.95$) or employment ($p = 0.81$) in the same year, ruling out confounding from local economic conditions.

The null is robust to alternative confirmation thresholds. Using 365 days (more inclusive) or 1,095 days (more conservative) produces reduced-form coefficients of 0.021 ($p = 0.81$) and 0.045 ($p = 0.59$), respectively—both near zero. There is no heterogeneity by baseline entry rates: the interaction between leniency and high-entry districts is insignificant.

6. Discussion

The null result—that randomly induced variation in bankruptcy debt relief does not increase business formation—has important implications for both theory and policy.

Why the null is informative. Three candidate mechanisms can explain why debt relief fails to promote entrepreneurship among Chapter 13 filers. First, the typical filer may lack the human capital, industry networks, or risk tolerance required for entrepreneurship, regardless of balance sheet position. Debt overhang may bind for investment by existing firms (Myers, 1977) without binding at the extensive margin of firm creation among consumer debtors. Second, the 3–5 year repayment period itself may delay or deter entrepreneurial plans: even with eventual confirmation, the years spent under court supervision may foreclose the optimal window for starting a business. Third, the credit market consequences of a bankruptcy filing—including the long-lived mark on the debtor's credit report—may offset the benefit of discharge by constraining access to startup capital (Keys et al., 2020).

Relation to prior work. The null is broadly consistent with Hurst and Lusardi (2004)'s finding that the wealth–entrepreneurship gradient is driven by preferences rather than liquidity constraints. It also complements Dobbie and Song (2015)'s finding that bankruptcy protection increases labor supply but does not increase self-employment income. Combined, these results paint a picture in which the fresh start facilitates recovery through wage employment—the margin most accessible to debtors—rather than through the more capital- and skill-intensive margin of entrepreneurship.

Limitations. Several caveats apply. The use of case duration as a proxy for confirmation introduces measurement error that attenuates estimates toward zero. The state-level outcome matching is imprecise for states with multiple bankruptcy districts. With only 10 court clusters, inference relies on asymptotic approximations that may be inexact. And the aggregation from individual cases to court-year panels sacrifices the case-level variation that makes judge-IV designs powerful. Future work linking individual bankruptcy records directly to state business registrations could provide sharper estimates at the individual margin.

7. Conclusion

The fresh start provided by Chapter 13 bankruptcy helps debtors recover—but it does not, by itself, create entrepreneurs. Using the cleanest available source of exogenous variation in consumer debt relief, this paper finds no detectable effect on business formation. The null is consistent across outcomes, time horizons, and specifications, and survives extensive diagnostic testing. The finding reframes the policy debate: the "fresh start" is valuable for what it demonstrably delivers—improved financial health, higher labor supply, better outcomes for children—rather than for the productive spillovers that debt overhang theory might predict. For the broader literature on financial constraints and entrepreneurship, the result underscores that the barriers to firm creation among consumer debtors lie beyond the balance sheet.

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

Angrist, Joshua D and Victor Lavy, "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement," *Quarterly Journal of Economics*, 1999, *114* (2), 533–575.

Bernstein, Shai, Emanuele Colonnelli, Xavier Giroud, and Benjamin Iverson, "Bankruptcy Spillovers," *Journal of Financial Economics*, 2019, *133* (3), 608–633.

Braucher, Jean, Dov Cohen, and Robert M Lawless, "Judges, Local Culture, and Bankruptcy Outcomes," *Emory Bankruptcy Developments Journal*, 2012, *29*, 291–370.

Chang, Tom and Antoinette Schoar, "Judge Specific Differences in Chapter 11 and Firm Outcomes," *NBER Working Paper*, 2016, (22632).

Dobbie, Will and Jae Song, "Debt Relief and Debtor Outcomes: Measuring the Effects of Consumer Bankruptcy Protection," *American Economic Review*, 2015, *105* (3), 1272–1311.

_ , Jacob Golán, and Crystal S Yang, "The Effects of Pre-Trial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges," *American Economic Review*, 2018, *108* (2), 201–240.

Evans, David S and Boyan Jovanovic, "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints," *Journal of Political Economy*, 1989, *97* (4), 808–827.

Frandsen, Brigham R, Lars J Lefgren, and Emily C Leslie, "Judging Judge Fixed Effects," *American Economic Review*, 2023, *113* (1), 253–277.

Haltiwanger, John C, "Business Applications as Economic Indicators," *Brookings Papers on Economic Activity*, 2021. Census Bureau Business Formation Statistics.

Hamdi, Nusret, Benjamin Iverson, and Karen Pence, "Intergenerational Effects of Consumer Bankruptcy Protection," *NBER Working Paper*, 2024, (32234).

Hombert, Johan, Antoinette Schoar, David Sraer, and David Thesmar, "Can Unemployment Insurance Spur Entrepreneurial Activity? Evidence from France," *Journal of Finance*, 2020, *75* (3), 1247–1285.

Hurst, Erik and Annamaria Lusardi, "Liquidity Constraints, Household Wealth, and Entrepreneurship," *Journal of Political Economy*, 2004, *112* (2), 319–347.

Keys, Benjamin J, Neale Mahoney, and Hanbin Yang, "Credit Market Consequences of Consumer Bankruptcy," *American Economic Review*, 2020, *110* (7), 2183–2219.

Kling, Jeffrey R, "Incarceration Length, Employment, and Earnings," *American Economic Review*, 2006, *96* (3), 863–876.

Maestas, Nicole, Kathleen J Mullen, and Alexander Strand, "Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt," *American Economic Review*, 2013, *103* (5), 1797–1829.

Myers, Stewart C, "Determinants of Corporate Borrowing," *Journal of Financial Economics*, 1977, *5* (2), 147–175.

Skeel, David A, *Debt's Dominion: A History of Bankruptcy Law in America*, Princeton University Press, 2001.

Sullivan, Teresa A, Elizabeth Warren, and Jay Lawrence Westbrook, "Less Stigma or More Financial Distress: An Empirical Analysis of the Extraordinary Increase in Bankruptcy Filings," *Stanford Law Review*, 2006, *59* (2), 213–256.

White, Michelle J, "Bankruptcy Reform and Credit Cards," *Journal of Economic Perspectives*, 2007, *21* (4), 175–199.

A. Data Appendix

CourtListener RECAP Archive. The RECAP Archive (https://www.courtlistener.com/) is a free, public repository of federal court documents maintained by the Free Law Project. I access case-level data through the REST API (v4), filtering for Chapter 13 bankruptcy cases (search type "r", query "chapter 13") filed between January 1, 2010 and December 31, 2019 in the 10 selected courts. For each case, I extract: court identifier, assigned judge name and ID, filing date, and termination date. I sample cases in quarterly batches (20 per court-quarter) and clean judge names by standardizing capitalization and removing system artifacts.

Census Business Dynamics Statistics. The BDS data are accessed through the Census Bureau's API (https://api.census.gov/data/timeseries/bds). I retrieve annual state-level counts of establishment entries (births), exits, total firms, total establishments, and employment for the 10 states containing the selected bankruptcy courts over 2010–2021. The BDS covers the universe of employer businesses with paid employees.

Confirmation proxy construction. I define plan confirmation using a duration threshold of 730 days. Cases with Duration > 730 are classified as confirmed; cases with Duration ≤ 730 are classified as dismissed or converted. This threshold is conservative: Chapter 13 plans run 36–60 months, so most confirmed plans will have durations well above 730 days. The threshold excludes early completions (rare in Chapter 13) and includes some cases that may have been dismissed after prolonged proceedings; both sources of misclassification bias estimates toward zero.

B. Identification Appendix

Balance tests. I regress pre-determined court-year characteristics (log case counts, number of active judges) on leave-one-out judge leniency with court and year fixed effects. Under random assignment, leniency should be orthogonal to these variables conditional on fixed effects.

Placebo test. I use contemporaneous (year t) business applications as a placebo outcome. Since Chapter 13 plans run 3–5 years, filers in year t are still in their repayment plans during year t and have not yet received discharge. Any effect of leniency on same-year business formation would indicate confounding rather than a causal channel through plan confirmation.

C. Robustness Appendix

Alternative thresholds. I re-estimate the main specification using 365-day and 1,095-day thresholds for the confirmation proxy. The 365-day threshold is more inclusive (classifying more cases as confirmed), while the 1,095-day threshold is more conservative.

Duration-based leniency. As an alternative to the binary confirmation proxy, I use the leave-one-out mean case duration for each judge as a continuous leniency measure. This avoids the arbitrary threshold choice but may be noisier.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
ln(Entry$_{t+1}$)	0.019	0.156	0.557	0.034	0.280	Small positive
ln(Entry$_{t+2}$)	-0.019	0.169	0.565	-0.033	0.300	Small negative
Entry rate$_{t+1}$	0.235	0.731	1.328	0.177	0.550	Large positive
ln(Firms$_{t+1}$)	-0.006	0.118	0.476	-0.013	0.247	Small negative
Panel B: Heterogeneous (by baseline entry rate)						
High-entry districts	-0.092	0.192	0.510	-0.180	0.376	Large negative
Low-entry districts	-0.036	0.267	0.276	-0.130	0.966	Moderate negative

Notes: **Country:** United States. **Research question:** Does consumer debt relief through Chapter 13 bankruptcy plan confirmation causally increase post-discharge business formation at the district level? **Policy mechanism:** Chapter 13 bankruptcy allows individual debtors to propose a 3–5 year repayment plan; judicial confirmation discharges remaining qualifying debts, relieving debt overhang that may constrain entrepreneurial entry. Judges vary in confirmation propensities within courts, creating exogenous variation in debt relief intensity. **Outcome definition:** Log annual new establishment entries from Census BDS (establishment births during the last 12 months), entry rate (entries as share of total establishments), and log total firms, at the state-year level. **Treatment:** Continuous—the share of Chapter 13 cases with confirmed plans in a court-year (instrumented by judge leniency). **Data:** CourtListener RECAP Archive for Chapter 13 cases (2010–2019, 10 federal bankruptcy courts, 6,016 cases, 94 judges) and Census BDS for establishment dynamics (state-year level, 2010–2021). **Method:** 2SLS using leave-one-out judge leniency as instrument for district confirmation rate; court and year fixed effects; standard errors clustered at court level. **Sample:** 10 large federal bankruptcy courts selected for high Chapter 13 volume; state-level outcomes matched to court jurisdictions. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the within-sample standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).