How do consumers fare when dealing with debt collectors? Evidence from out-of-court settlements

Approximately 14 percent of U.S. consumers experience debt collection each year. Policymakers have long worried that many of these consumers may be unaware of their rights when negotiating with debt collectors, making them vulnerable to being taken advantage of. However, there is virtually no evidence on whether these concerns are well founded. In this paper, we investigate whether consumers who negotiate directly with debt collectors agree to deals that suggest they are ill-informed. To do so, we exploit a setting where randomly assigned civil court judges have different propensities to encourage consumers and collectors to negotiate with one another before a trial. Using new data that links court records from debt collection cases with subsequent consumer credit outcomes, we find that consumers who settle causally experience higher short-term financial distress than if their case had been resolved in court. Overall, the evidence suggests that consumers fail to walk away from bad deals offered by debt collectors.

JEL codes: D14, D18, G00

Keywords: Debt collection, negotiation, financial settlements

How do consumers fare when dealing with debt collectors? This question is one of growing importance. Approximately 14 percent of U.S. consumers have been under third-party debt collection in recent years (Federal Reserve Bank of New York, 2015), and the industry collects over $55 billion annually (Ernst and Young, 2013). Typically, debt collectors attempt to negotiate payments directly with consumers, or use the court system to extract payments. Policymakers have long worried that, when collectors negotiate directly with consumers, they may engage in abusive or deceptive practices, or that consumers are unaware of their rights. Among other things, collectors may take advantage of consumers who do not know what will happen if they fail to pay. Reflecting these concerns, provisions of the Fair Debt Collection Practices Act (FDCPA) and Dodd-Frank Act seek to prevent debt collectors from misleading consumers and require them to provide fair notice of a consumer's obligations. The Consumer Financial Protection Bureau (CFPB), the regulator charged with enforcing the relevant provisions of these laws, provides several resources that seek to educate consumers about their rights when dealing with collectors.

However, there is virtually no evidence on whether these efforts have been sufficient to protect consumers, nor whether the concerns motivating these efforts were well-founded in the first place. This paper investigates whether consumers who negotiate directly with debt collectors agree to bad deals that suggest they are ill-informed about their alternative options.

A bad deal is one that makes a negotiating party worse off than they otherwise would be if they rejected the deal and exercised their outside option. A rational agent with full information should never agree to a bad deal. In the context of debt collection, consumers' outside options are largely determined by the legal system, which has various consumer protections built into it. If consumers are ill-informed about these options, then making a deal could cause them to be worse off than they would be otherwise. Following this logic, we examine whether consumers agree to bad deals by examining whether making deals has a negative causal effect on their outcomes.

There are two major challenges in estimating the causal effect of making a deal with a debt collector. The first is measurement: it is difficult to observe negotiated agreements and subsequent consumer outcomes. The second challenge is identification: negotiated agreements do not occur randomly and depend on unobservable consumer characteristics. For example, consumers who make deals may tend to be unobservably wealthier than others. In this case, even if making a deal has a negative causal effect on these consumers, they may have better outcomes than observationally similar consumers who do not make a deal.

To address these challenges, we conduct our study in a controlled setting. Specifically, we study consumers who face litigation by a debt collector in state civil court. Litigation is one of the primary tools that collectors use to extract payments from consumers. In such lawsuits, a collector seeks a court judgment certifying the legal validity of their claim. Unlike federal bankruptcy court, the court is not a forum to discharge or consolidate debt. Rather, the role of the court and presiding judge, is to verify that the collector has the proper legal status. If the collector wins (as they mostly do), they are then entitled to, among other things, garnish the borrower's wages or bank accounts up to statutory limits; if they lose, the status quo resumes.

In this setting, the collector and consumer may reach an out-of-court negotiated settlement such as a lump-sum payment in lieu of having the case heard. Crucially, borrowers may reject a proposed settlement and exercise their outside option to have the court hear their case. In our setting, this outside option is very clear and is the same for all consumers in our sample. Our main finding is that consumers who settle causally experience significantly higher short-term financial distress than they would have compared to this outside option, and that these are likely bad deals.

We exploit several features of our setting to reach this conclusion. To address the measurement challenge, we assemble a unique dataset that links all court records from Missouri debt collection lawsuits from 2007-2014 with credit registry data from TransUnion. From the court records, we observe cases that concluded with an in-court ruling as well as cases that concluded with an out-of-court settlement between the two parties. The credit registry data allow us to track subsequent financial outcomes for each consumer. We focus on Missouri because, unlike most states, it has a centralized database of cases tried in different circuit courts. Importantly, this database also tracks garnishments. Missouri is also a representative state in terms of collection (Ratcliffe et al., 2014).

To address identification, we exploit the fact that judges are randomly assigned and have different empirical propensities to preside over cases that settle out of court. We attribute these different settlement propensities to differences in the style of how judges manage their case docket.

According to Missouri debt collection attorneys we spoke to, prior to hearing a case, there is variation across judges in how much they encourage the parties to reach a settlement. For example, at the start of the day, some judges routinely ask all parties to talk to one another first to try to reach an agreement, while others do not. Consistent with the idea that settlement propensity is related to a judge's style, we find that is it persistent over time for a given judge. Because judges are assigned randomly, we can then instrument for settlement with judge settlement propensity. While settlement may be correlated with unobservable consumer characteristics that relate to financial distress (e.g. wealth), judge settlement propensity should not be.

Of course, judges may be able to take multiple actions with the potential to influence the subsequent financial distress of consumers. For example, it may be that high settlement propensity judges do tend to encourage out-of-court settlements by prompting negotiations but that they also tend to rule against consumers in cases that go to trial. If so, it is possible that randomly drawing such a judge does cause an increase the probability of financial distress, on average, but only though the negative-ruling channel, not through the prompting-negotiations channel. To address this possibility, we control directly for judges' tendency to rule against consumers throughout our analysis. Thus, our identifying variation stems from variation in judges' settlement propensities that is orthogonal to judges' ruling tendencies.

We begin by hand-verifying that judges in our sample are randomly assigned by obtaining the court procedure documents and speaking with the court clerk for every court district in Missouri. We exclude those (mostly low-population) counties that do not use random assignment. In our sample of randomly assigned districts, we find no significant differences in the year prior to case disposition among borrowers who draw a higher settlement-prone judge versus those who draw a lower settlement-prone judge, consistent with random assignment. These borrowers have very similar credit scores, revolving balances, mortgage balances, and other characteristics.

After case disposition, differences emerge. More precisely, when we instrument for settlement using the settlement-propensity of the judge, we find that settlement causes an increase in short-term household financial distress by significantly increasing the probability of delinquency, bankruptcy, and foreclosure in the first year after case disposition. The effects are economically meaningful. Settlement triples the probability of bankruptcy and doubles the probability of foreclosure, with effects concentrated in the first year after case disposition. The results are robust across various specifications and sample restrictions. For example, we find similar results when we exclude default judgments. Consistent with what we would expect given judges' limited discretion, our results are virtually the same with and without the control for in-court leniency. Household financial distress likely arises due to liquidity needs, as we find a significant increase in mortgage balances along with the increased risk of foreclosure.

Are these bad deals? It is possible that making a deal causes consumers to be worse off in terms of financial distress, but better off along some other dimension. Taking such tradeoffs into account, consumers may fully understand their options, but prefer making a deal. For this to be the case, there must be benefits to making a deal. One such potential benefit might be preserving access to credit, as losing at trial could adversely affect a consumer's credit score. However, the population we study in this paper already have significant negative flags on their credit reports, including severe delinquencies and accounts in collection. Therefore, it is unlikely that the benefit of avoiding one more potential negative flag from a court judgement is that great. Consistent with this idea, we find that the average credit score of consumers in our sample prior their case is only 536, and that the effect of settlement on a consumer's credit score is statistically indistinguishable from zero. Moreover, we find that the effect of settlement on financial distress does not vary with a consumer's pre-case credit score, even though consumers with high credit scores should be more willing to strike deals that risk financial distress.

Another potential benefit from making a deal is that consumers may be able to extract concessions from collectors who want to avoid the costs of going to trial and collecting payments through garnishment. Garnishment involves an implicit haircut on average, as collectors may end up garnishing less than the full amount if the borrower declares bankruptcy, loses their job, or for other reasons. Data on garnishment cash flows shows that the average garnishment haircut is 62% assuming a discount rate of zero. Our data do not include private settlement amounts, but consumers who settle would need to obtain a haircut of at least this amount to break even with the expected payouts from going to court. To justify a good deal and compensate for the higher risk of financial distress, consumers would need to extract even larger haircuts, up to a maximum of 62% plus the collector's court cost. We view this as possible but unlikely: settlements likely involve non-trivial amounts given the subsequent financial distress, and professional collectors face arguably low marginal court costs.

Finally, another potential benefit from making a deal is that it reduces uncertainty. Consumers may be risk-averse and thus willing to bear higher expected financial distress from settlement to avoid the gamble of going to court. However, most consumers lose in court; therefore, they do not actually face that much uncertainty from choosing to go to trial. Moreover, we find that consumers who settle actually experience more financial distress on average than both trial winners and trial losers. Thus, settling leads to an outcome that is worse that than either the good or bad outcome from going to court. No level of risk aversion would justify choosing a dominated outcome distribution of this kind.

How likely are our findings to apply beyond our specific setting? As stated earlier, approximately 14% of consumers with a credit file have at least one account in third-party collection each year. These consumers routinely negotiate deals with collectors, often without being litigated. While we only study those who are litigated, it seems plausible that those who are not may similarly make bad deals. In fact, one could argue that bad deals that trigger financial distress may be even more prevalent among those who are not litigated, as those who are litigated are likely selected by collectors based their ability to pay. Moreover, the types of debt collectors who operate through the court system and negotiate at the behest of a judge likely engage in fewer nefarious practices than others, leading to fewer bad deals. However, our findings suggest that consumers even make bad deals under these circumstances.

On the other hand, since we employ an instrumental variables strategy, we can only identify a local average treatment effect (LATE), or the causal effect of settlement among those for whom drawing a high settlement-propensity judge is pivotal in the settlement decision. Therefore, one possible concern is that settlements may on average be good for consumers, but the settlements that judges specifically induce are bad. For that to be the case, it would have to be that consumers are actually well-informed but that drawing a high settlement-propensity judge causes them to become ill-informed, perhaps because they misinterpret the judge's prompt to negotiate as a signal that they are less protected under the legal system than they thought. However, if consumers' beliefs about their rights are so easily swayed, they were likely never actually that well informed to begin with and could have been convinced by a debt collector to make a bad deal even without a judge's prompt to negotiate.

The main contribution of this paper is to provide new evidence that a large, previously understudied group of consumers are, on average, making bad deals that are highly favorable to collectors. To date, we only have scant evidence of about how consumers interact with debt collectors (Hunt, 2007; Zinman, 2015), and our paper thus directly contributes to literature on the role of consumer protections and decision-making. While policymakers have focused on potentially abusive practices by collectors, the result that consumers make bad deals even when collectors follow the law and litigate in court suggests that there is potentially an important on-going role for financial education with respect to debt collection.

This paper also connects to the broader literatures on household financial distress and debt relief. Previous research has shown that the consumers who declare bankruptcy are typically very financially distressed. Relief from this financial distress via personal bankruptcy can improve earnings, mortality and financial health (Dobbie and Song 2016, Dobbie et al. 2017) and can have other real effects (Musto, 2005, Dobbie et al. 2016, Herkenhoff et al., 2016 Severino and Brown, 2017). Our results suggest that making a deal with a collector increases the probability of a consumer becoming financially distressed to the point where she must seek debt relief.

1. Hypothesis development

1.1. Background

Between 2007 and 2014, on average 6.0% of total consumer debt, or about 711 billion dollars, has been 90 days delinquent or later.[1] Consumers who fall significantly behind on their debt

[1] This discussion combines observations from Hunt (2007) as well as updated statistics and author calculations from the Federal Reserve Bank of New York (2018)'s quarterly report on household debt and credit. To calculate the number of consumers with accounts in third party collection and the fractional flow rate of bankruptcy, we assumed there were 200 million consumers with a credit file, a conservative number (Lee and van der Klaauw, 2010).

payments enter the collections process. Several papers have focused on the foreclosure crisis during the Great Recession, but the literature has paid comparatively less attention to consumers who have been subject to collections more broadly. While 1.4 million consumers experienced foreclosure each year on average, a much larger number—14% of consumers with a credit file, or around 33 million—have at least one account in third-party collection each year, with an average balance of 1,400 dollars. In aggregate, relatively few of these consumers end up in bankruptcy (Dawsey and Ausubel, 2004; Dawsey, Hynes, and Ausubel, 2013; Fay, Hurst, and White, 2002; White, 1998)—the average flow rate into bankruptcy each year is 1.6 million consumers.

For short-term delinquencies, lenders often rely on in-house collections departments. For severely delinquent debt, many lenders rely on third-party collection agencies working on a fee-basis or sell off the debt outright to debt buyers. The key challenges facing collectors are to first determine and then locate which consumers likely have the financial resources to pay, and then to extract as much as cash as possible before other agents do for potentially other delinquent accounts. For unsecured debt, their tools are negotiation and litigation. Litigation often leads to wage garnishment, the extraction of cash flows from a consumer's paycheck (Hynes 2005, 2008).

Policymakers have had long-standing concerns potentially abusive collection practices and the possibility that consumers do not fare well when making deals with collectors. Several consumer protection laws and regulations circumscribe how lenders or third-party collectors may interact with consumers. Among other protections, the Fair Debt Collection Practices Act (FDCPA) of 1977 limits when third-party collectors can contact a consumer, prohibits misrepresentation, lies, and deception, and prohibits the collection of amounts greater than the amount owed, which the collector must provide written notification of in a timely fashion. Several individual states have enacted longstanding laws that provide protections that go beyond the FDCPA (Fedaseyeu, 2015). More recently, the Dodd-Frank Act of 2010 empowered the newly-created CFPB to oversee the industry. Motivated by their observation that "debt collection constitutes one of today's most important consumer financial concerns" (CFPB, 2014), the CFPB has been exploring potential new rules for the industry and publishes several educational resources for consumers. Despite these longstanding regulations, however, concerns persist, largely because a lack of data has limited our understanding of how consumers fare when dealing with collectors (Zinman, 2015).

1.2. Collection in Missouri civil courts

We study how consumers and collectors fare in negotiations that occur after a collector has sued a consumer in state civil court to collect on non-mortgage debt, but before the court resolves the case. We further limit our study to Missouri because, unlike most states, it has a centralized database of cases tried in different civil courts. Missouri is a representative state in terms of percentage of consumers who are delinquent and the average amount of debt in collections (Ratcliffe et al., 2014). Missouri is also not particularly exceptional with regards to the law surrounding collections; a few other states such as Texas and Pennsylvania severely curtail the ability of collectors to garnish wages. Focusing the scope of our study on this setting affords us several advantages in testing our main hypothesis, as we discuss in Section 2.

Before proceeding, review the key institutional details surrounding the litigation process. The plaintiff in a case is typically an attorney acting on behalf of the original lender or a debt buyer and seeks a court judgment certifying the validity of the debt. Debtholders can sue at any point before the state's statute of limitations expires on the debt, which is 10 years in Missouri. After the statute of limitations expires, the debt is still owed, but the debt holder legally cannot use the court system to aid in the collection of the debt.

1.2.1. Timing of case resolution. After the debt holder files a lawsuit, the court attempts to serve the borrower with a summons. If the borrower cannot be located, the case is dismissed without prejudice, meaning that the debt holder does not win the case but retains the right bring the case again in the future. If the borrower is successfully served, he must appear in court on the assigned date. If the borrower fails to appear on this "first appearance" date, the debt holder typically wins a "default judgment" against the borrower. If the borrower does appear, a subsequent hearing date is set, before which there may be additional court appearances required if additional legal issues come to the fore.

At the hearing, the judge determines whether there is sufficient evidence that the plaintiff has the proper legal status and documentation to collect on the debt. If the consumer does happen to win, the court dismisses the case, typically "without prejudice," meaning the plaintiff may bring another case in the future (perhaps with the right documentation). In rarer instances, the court dismisses the case "with prejudice" if the judge has determined that the plaintiff should not be able to bring another lawsuit against the consumer, for example if the plaintiff has acted improperly.

If the plaintiff wins—the usual outcome—the court enters a judgment against the borrower. These often take the form of "Consent Judgments," essentially court-sanctioned payment plans where a consumer has admitted they owe the debt to the collector. Typically, the judgment amount is for the principal plus interest and court fees. A judgment grants the debt holder the right to garnish the borrower's wages or bank accounts up to certain statutory limits. Missouri's cap wage garnishment is 10% of disposable income, which is more stringent than the federal limit of 25%. The judgment itself does not initiate garnishment; the debt holder must separately contact the consumer's employer and notify them that they have a court-sanctioned judgment that allows them to garnish wages before receiving cash flows.

At any point in this process, a borrower can reach a negotiated settlement agreement with the collector. If they reach a settlement before a hearing, the court closes the case and records the case as having been dismissed by the parties involved. These settlements are often for a lump-sum payment, as even if consumers agree to a payment plan, collectors will often go to the court to obtain a judgment as the consumer may fail to pay. As mentioned earlier, there are typically several court appearances before an actual trial, and thus several opportunities for negotiation. Once a case has been settled or heard in a court, the court records the case as "disposed."

1.2.2. The role of judges. The judge determines the legal validity of the collector's claim. Importantly, the court is not a venue for debt discharge or consolidation, and the scope of items for the court to determine is relatively narrow. For example, according to conversations with Missouri debt collection attorneys, consumers often provide non-legal arguments as to why they should win that are not relevant from a legal perspective and that the courts routinely disregard.

In contrast, our conversations with Missouri attorneys suggest that judges can and do influence whether parties make deals on these appearance dates to varying degrees as part of how they manage their daily case docket. For example, on the first appearance date, the judge may state that if the two sides have not talked they should see if something can be worked out. Typically, the case is set for trial at a later appearance. Before scheduling a trial date, some judges will again encourage the two sides to negotiate. Finally, on the trial date, some judges will one last time encourage talks, as by this point there is sometimes new information to discuss. Courts often have tables or side rooms designated for such negotiations. As noted above, many borrowers do not have a lawyer and rely on non-legal arguments, which are effectively a waste of the court's time.

1.3. Conceptual framework

We evaluate how consumers fare in negotiations by testing whether consumers follow a fundamental principle of negotiation when dealing with collectors: one should reject bad deals. A "bad deal" is a proposed settlement that makes one worse off than the one's outside option.

Suppose a collector and consumer negotiate over a lump-sum payment of s dollars. If negotiation breaks down, the consumer and collector resolve the case in court, where the collector wins judgment J (the present value of expected cash flows associated with garnishment, or another court-sanctioned payment plan) with probability p. If the consumer settles and makes a lump-sum payment, she incurs financial distress Z. If she goes to court and loses, she experiences financial distress X from wage garnishment. Because payment methods are different, Z may not equal X.[2] Payoffs are then:

	Consumer	Collector
Settle	$-(s+Z)$	s
Court	$-p(J+X)$	pJ

In this setup, a consumer should never take a "bad deal" where $s > pJ+(pX\text{-}Z)$, because such a deal has $s+Z > p(J+X)$.

Any equilibrium with settlement excludes bad deals. "No bad deals" requires consumers only agree to deals with weakly positive consumer surplus, $0 \leq (pJ\text{-}s)+(pX\text{-}Z)$. On the other hand, "no

[2] One should interpret Z and X as incremental to the financial distress the consumer is experiencing immediately before she enters the negotiation. This allows us to write the payoff of the consumer winning in court as normalized to zero, as we assume that the consumer does not pay any incremental cash flows if they win.

bad deals" also requires collectors only accept deals with $s \geq pJ$. Thus, in any equilibrium with settlement, we must have $pX\text{-}Z \geq 0$, irrespective of the bargaining power of each party and the equilibrium settlement amount s. Bargaining power affects how the parties divide any available total surplus, but if $pX\text{-}Z$ is negative, the total surplus is negative, and no deal is acceptable to both parties. We assume that collectors act rationally, so that, if we observe $pX\text{-}Z < 0$, we interpret this as evidence that consumers taking a bad deal. While the framework above fixes distress costs Z and X for simplicity, a similar insight holds if they scale with s and J.

Prediction 1 (No bad deals). *In any settlement equilibrium, we must have $Z \leq pX$: Financial distress incurred by a settlement should be weakly less than that incurred by going through court.*

We focus our main empirical analysis on testing Prediction 1. We focus on the probability of subsequent delinquency, bankruptcy, and foreclosure as our dependent variables. Our choices of bankruptcy and foreclosure are important in that they provide strong tests of whether consumers are striking bad deals with $Z > pX$. If we find that settlement causes such high distress that consumers significantly increase their risk of bankruptcy and foreclosure, this is strong evidence that $Z\text{-}pX$ is positive and large and that consumers are making bad deals.

We consider three possible alternatives that may make higher financial distress part of a good deal. First, one alternative is that, despite the increased bankruptcy and foreclosure risk, these are good deals for consumers because the deals avoid damaging their credit reports with judgments. In terms of the framework, this alternative is that pX is in fact quite large due when including this credit report damage, so that $Z \leq pX$ despite the higher bankruptcy and foreclosure risk. We address this possibility by examining how much settlement improves credit scores and by testing our hypothesis within a subset of consumers whose credit reports are effectively beyond repair. For these consumers, the marginal damage to their credit report from going to court should be small and should contribute little to pX.

Prediction 2 (No bad deals with credit score improvement). *An equilibrium settlement may involve financial distress with higher delinquency, bankruptcy, or foreclosure risk compared with going to court if $Z \leq pX$ where pX includes the damage to a consumer's credit score.*

Second, and more broadly, one might worry that higher financial distress $Z > pX$ could be part of a good deal for the consumer if other sources of surplus offset this cost. Suppose the collector and consumer earn generic surpluses of F and G, respectively, if they settle and avoid court. The consumer's no-bad-deal condition is then $0 \leq (pJ\text{-}s)+(pX\text{-}Z)+G$. For higher financial distress $Z > pX$ to be part of a good deal, the consumer benefit G plus any concessions by collectors $pJ\text{-}s$ must exceed the cost of this higher distress. Concessions ($pJ\text{-}s > 0$) can arise if consumers have significant bargaining power, as collectors may be willing to concede up to F due to their own no-bad-deals condition: $pJ\text{-}s \leq F$. For small G, such concessions from the collector are a necessary condition for a good deal.

One obvious candidate for F and G are legal fees. However, most consumers who show up to court do not have a lawyer, so G is likely small for the court outcomes that we see, or at least potentially smaller than $Z\text{-}pX$. Finding that $pJ\text{-}s \leq 0$ and $Z > pX$ then implies consumers who settle both pay more and experience higher financial distress than going to court without a lawyer, providing even stronger evidence of a bad deal.

Prediction 3 (No bad deals with other sources of surplus). *Any equilibrium settlement with $Z > pX$ and $G < Z\text{-}pX$ must involve concessions from collectors, $pJ\text{-}s > 0$.*

Third, risk aversion may also make higher financial distress $Z > pX$ part of a good deal. Consumers may be willing to endure higher financial distress as part of the certainty-equivalent of avoiding the gamble of going to court. In utility terms, the no-bad-deal condition for a consumer is $u(-(s+Z)) \geq u(-p(J+X))$, where u is the consumer's utility function. We argue this is unlikely since most consumers lose in court and p is high. Setting this aside, however, note that, if u is increasing and $p < 1$, we must have $u(-(J+X)) < u(-p(J+X))$. It follows that sufficient conditions for a bad deal are $Z \geq X$ and $s \geq J$ with strict inequality for at least one, because then $u(-(s+Z)) < u(-(J+X)) < u(-p(J+X))$. Intuitively, risk aversion cannot explain why consumers who settle end up worse off than those who lose in court.

Prediction 4 (Sufficient conditions for a bad deal). *If $Z \geq X$ and $s \geq J$ with strict inequality for at least one, then, without other surplus, risk-averse consumers made a bad deal.*

2. Data and variables

We assemble a unique dataset that contains all court records from Missouri debt collection lawsuits from 2007-2014 merged with credit registry data from TransUnion.

Our starting point for the data are the court records. To begin with, there are 667,337 debt collection cases in Missouri's 45 court districts during our sample period. Our empirical design ultimately focuses on court districts where we are able to verify random judge assignment. We first obtain and hand-review the court procedure documents and look for evidence of random assignment. We then proceed to call the district court clerk on two separate occasions and speak with them about their practices to verify random assignment. This leaves us 203,298 cases in 10 court districts. Court districts correspond to one or more counties. The counties corresponding to the districts in our sample include high population counties (e.g., Jackson County) and typically exclude low population ones (e.g., Ozark County), where there may only be one relevant judge. Some high population counties (e.g., St. Louis) do not assign cases completely randomly, so we exclude them.

To examine consumer credit outcomes after case disposition, we link the court records with detailed credit registry data from TransUnion. This link was performed by TransUnion based on names and standardized addresses as well as birthdates and social security numbers when available. We purchase 9 years of credit files, 2007-2015, to match with the court data. Each credit file contains a snapshot of the consumer's credit profile in the January of that year. To preserve anonymity, the matched data returned to us by TransUnion was stripped of these identifiers. TransUnion was able to match approximately 87% of consumers from the court records to their database leaving us with a sample of 176,769 cases.[3]

[3] This match rate is somewhat higher than that of previous research. One possible reason is that, rather than attempting to match a consumer separately each year, we asked TransUnion to pull a consumer's file for all years if the name/address we provide matches in any year.

After applying three further filters, our final sample consists of 82,218 cases heard by 43 judges. First, we require that the consumer's case is heard by a judge for which we can construct a judge settlement-propensity measure described in Section 2. Second, we require that we observe the consumer's credit file both in the January before disposition (which we will call time 0) and in the January after. Finally, we require that the data indicates the borrower was appropriately served.

From the court records, we can observe cases that concluded with a ruling in favor of one party or the other. Following the guidance of the court clerks we spoke to, we categorize a case a being settled out of court if the defendant was successfully served but the case was ultimately "Dismissed by Parties." Figure 1 shows the type of case outcomes included in the analysis: "Settlement" refers to an out-of-court bilateral arrangement and represent about 17% of the outcomes. "Consent judgment" represents 17% of outcomes, dismissal (with or without prejudice) represents 5%, and "default judgment" represents 62%.

Table 1 Panel A reports summary statistics. Conditional on no settlement, 94% of cases end in judgment, and 6% end in dismissal, with an average judgment amount of 4,149 dollars. Only 50% of these judgments lead to subsequent garnishment cash flows. The average length of time from filing to disposition is 88 days. Most cases in our sample correspond to a unique defendant.

Table 1 Panel B reports the characteristics of borrowers in in our sample in the year before the disposition date of a case (column 1) and compares these with the overall population of credit users (column 2) as well as the population of borrowers who declare bankruptcy (columns 3 and 4). They have comparable age and slightly lower homeownership rates than bankruptcy filers. The average credit score of consumers in our sample is quite low, 536, and is lower than both the average population score and the score of bankruptcy filers. Not surprisingly, they have a higher likelihood of having a collection flag (76% vs. 47%), with higher collection balances.[4]

3. Empirical strategy

3.1. Empirical specification

[4] This fraction is less than 1 because collectors do not always report collections on a consumer to the credit bureaus.

We are interested in estimating the effect of the out-of-court settlement on consumer financial distress after a case is disposed. A naïve empirical design would the following using OLS:

$$y_i = \alpha_{cs} + \beta S_i + \Gamma_0 X_i + \Gamma_1 J_{ijcs} + u_i, \tag{1}$$

where y_i is consumer i's outcome in the period of interest (e.g., 1 year after disposition), S_i is an indicator of whether the case was settled out of court, α_{cs} is a court c-by-calendar disposition year s fixed effect (e.g., Jackson Circuit x 2008) to account for court specific time-varying trends, and X_i is a set of controls that include age, credit score, days-to-disposition, homeownership status, and a flag for previous bankruptcy filings, measured before in the January before case disposition (denoted as year 0). We bin age (5-year bins), credit score (50 point bins), and days-to-disposition (30 day bins). The variable J_{ijcs} reflects the tendency for a judge j assigned to case i to rule for or against a consumer, a control variable we discuss in more detail below.

Our three primary credit variables of interest y_i are flags for whether a consumer is delinquent on debt, has recently filed for bankruptcy, and has recently experienced foreclosure. However, the error term u_i in Equation 1 likely contains unobserved borrower characteristics affecting consumer financial distress that are correlated with settlement S_i: $E\left[S_i u_i | X_i, J_{ijcs}\right] \neq 0$. This biases estimates of the effect of out-of-court settlements even with the inclusion of court-disposition year fixed effects. Underlining this concern, Table 2 suggests that consumers who settle statistically and differ from those who do not in several observable dimensions. Consumers who settle tend to have higher average credit scores (562 vs. 530), mortgage balances ($46- vs. $28-thousand), more trade lines (4 vs. 3), and lower collection balances ($5800 vs. $7300). This suggests that they may be unobservably wealthier and less likely to file for bankruptcy or experience foreclosure than consumers who do not settle.

To overcome this identification challenge, our empirical strategy exploits the random assignment of judges. There is significant variation in the fraction of cases a judge presides over that end with a settlement, or a judge's "settlement-propensity." This is consistent with variation in how judges manage their case dockets and in how much they encourage parties to settle, as discussed in Section 1.2.

Considering these differences in judge style, we estimate judge-year specific settlement propensities following a leave-out estimate methodology. Specifically, we compute:

Judge settlement propensity: $SP_{ijct} = \frac{\sum_{k=1}^{n_{jct}} s_k - s_i}{n_{jct}-1} - \frac{\sum_{k=1}^{n_{ct}} s_k - s_i}{n_{ct}-1}$, (2)

where n_{jct} is the number of cases judge *j* in court *c* hears in year *t* and n_{ct} is the number of cases heard by the broader court. This ratio SP_{ijct} represents the leave-out average settlement rate of judge *j* in court *c* in year *t* minus the rate in court *c* in year *t* (see, e.g., Kling 2006; Chang and Schoar 2008; Doyle 2007, 2008; Aizer and Doyle, 2013; Dobbie and Song, 2015; and Dobbie et al. 2017). We follow Dobbie and Song (2015) in subtracting the leave-out average settlement rate of the broader court to remove any court-level heterogeneity in settlement rates. We first estimate judge settlement propensities in the full unmatched sample of cases and include only cases where the judge heard a minimum of 10 cases and had a 5% case share per judge-year within the final sample of cases where we confirmed the defendant was served.

We use judge settlement-propensity SP_{ijcs} to instrument for settlement S_i in Equation 1. Specifically, the first-stage equation is:

First stage: $S_i = a_{cs} + b\, SP_{ijcs} + \mathrm{G}_0 X_i + G_1\, J_{ijcs} + v_i$, (3)

where SP_{ijcs} is the leave-out settlement rate in Equation 2 and the remaining variables are defined as in Equation 1. The second stage equation is Equation 1, where we estimate the parameters using standard instrumental variable techniques:

Second stage: $y_i = \alpha_{cs} + \beta S_i + \Gamma_0 X_i + \Gamma_1\, J_{ijcs} + u_i$. (4)

We cluster standard errors at the judge level to account for across time correlations between cases and cross-sectional co-movements within a judge-court-year.[5]

Our identifying assumption is that the random assignment of judge with different settlement propensities generates variation in the probability that the two parties settle that is orthogonal to consumer heterogeneity: $E\left[SP_{ijcs} u_i | X_i, J_{ijcs}\right] = 0$. Under this assumption, the second-stage coefficient β on settlement is the causal impact of settlement on subsequent outcomes relative to

[5] We confirm that results are similar if we cluster at the court level while correcting for small-cluster bias.

walking away from the negotiating table. Under the null hypothesis that consumers reject bad deals, this β coefficient should weakly indicate less financial distress.

If the identification assumption holds, consumers assigned a higher settlement-propensity judge have similar characteristics as consumers assigned a lower settlement-propensity judge yet are more likely to conclude their case with a settlement, while nonetheless keeping the same outside option to have their case heard in court. In reduced form, our empirical strategy compares the financial outcomes of these two groups of consumers.

Of course, judges may be able to take multiple actions with the potential to influence the subsequent financial distress of consumers. For example, it may be that high settlement propensity judges do tend to encourage out-of-court settlements by prompting negotiations but that they also tend to rule against consumers in cases that go to trial. This motivates our inclusion of J_{ijcs} in Equations 3 and 4: our identifying variation comes from the component of judge settlement propensity orthogonal to a judge's ruling tendencies in-court.

While the identification assumption is untestable, we nonetheless can test whether SP_{ijcs} is correlated with observable characteristics. If judges are randomly assigned, it should uncorrelated. We sort cases by whether SP_{ijcs} is above or below the sample median and report the average year-0 consumer characteristics across these two groups in Table 2 Panel B. Unlike Panel A, we find no significant economic or statistical differences in the year prior to case disposition among borrowers who draw a higher settlement-propensity judge versus those who draw a lower settlement- propensity judge, consistent with random assignment.

4. Prediction 1: Bad deals and financial distress

4.1. First stage

We start by graphically depicting the key element of our first-stage relationship in Figure 2. The figure plots a settlement indicator against our leave-one-out measure of judge settlement propensity. To construct the binned scatter plot, we first regress an indicator for settlement on court-by-year fixed effects and calculate residuals. We then take the mean residual in each judge-by-year bin, adding the mean settlement rate to each residual to aid in the interpretation of the plot. The solid line shows the best linear fit estimated on the underlying microdata estimated using OLS.

Table 3 reports the first stage in regression form. The first column corresponds to Figure 2. The additional columns add additional controls. As can be seen, we estimate a strong positive relation between judge settlement propensity and settlement. The F-statistic, shown below the estimated coefficients, is quite high and easily surpasses the rule-of-thumb threshold for weak instruments (Stock and Yogo, 2005).

4.2 IV estimation

Table 4 Panel A reports our main result from estimating Equation 4 for financial outcomes the year after case diposition. Columns 1, 3, and 5 report results from estimating Equation 4 using OLS for the dependent variables of delinquency, bankruptcy, and foreclosure. The coefficients suggest that borrowers that settle are 5% less likely to file for bankruptcy and 1% less likely to experience foreclosure a year after case disposition, with effectively zero effect on delinquency rates. Given the average bankruptcy and foreclosure rates of 5% and 3% among borrowers who go court (Table 2, Column 2), these are potentially large effects. However, this could be driven by unobserved differences between borrowers who settle and those who do not. For example, borrowers who settle have be wealthier and less likely to declare bankruptcy, as Table 2 suggests.

Columns 2, 4, and 6 of Table 4 Panel A reports results from the instrumental variables estimation that uses the judge settlement-propensity SP_{ijct} to instrument for settlement S_i. Settlement leads to 13%, 11%, and 4% increases in the probability of delinquency, bankruptcy, and foreclosure, which are statistically reliably different from zero at the 10%, 1%, and 1% levels, respectively. The effects are economically significant as well. Relative to the base rates of delinquency, bankruptcy, and foreclosure reported in Table 2 (columns 4 and 5) of 53%, 7%, and 3%, the point estimates in Table 4 suggest that settlement increases these rates by a multiple of 1.2, 2.6, and 2.3, respectively.

Given our identification strategy, we interpret these as causal local average treatment effects: individuals who were induced to settle through the random judge they drew would have on average experienced lower rates of delinquency, bankruptcy, and foreclosure had they not settled. Comparing the OLS estimates and IV estimates in Panel A highlights the severity of the endogeneity problem when attempting to isolate this causal effect. Unobserved consumer heterogeneity such as consumer wealth confound the endogenous OLS correlations between settlement and financial distress.

Overall, the results reject Prediction 1: considering only these variables, consumers who settle make deals with $Z > pX$ and incur more financial distress, leading to higher rates of subsequent bankruptcy and foreclosure. Table 4 Panel B explores how this distress unfolds by estimating Equation 4 with dollar credit balances as left-hand side variables. Settlement causes borrowers' mortgage balances to increase by a point estimate of $30,000 (reliably different from zero at the 1% level), potentially to finance their settlement or post-settlement expenses. This is consistent with the higher rates of foreclosure in Panel A. Point estimates suggest that settlement leads borrowers to increase their revolving balance by $2,500, but this estimate is not statistically reliably different from zero. Note that the increase in credit balances does not necessarily indicate that consumers who settle are accessing more credit. It may be that they are simply paying down their debt more slowly (e.g. by missing mortgage payments).

4.3 Robustness

4.3.1. What about other judge actions? The random assignment of judges ensures that we are estimating a causal effect stemming from judge actions, rather than a correlation driven by unobservable borrower characteristics. However, judges may be able to take multiple actions with the potential to influence the subsequent financial distress of defendants, and their tendency to take these actions may be correlated. For example, it may be that high settlement propensity judges do tend to encourage out-of-court settlements by prompting negotiations but that they also tend to rule against consumers in cases that go to trial. If so, it is possible that randomly drawing such a judge does cause an increase the probability of financial distress, on average, but only though the negative-ruling channel, not through the prompting-negotiations channel. In particular, it may be that the consumers who go to trial are worse off as a result of drawing such judges, but those who settle are no worse off. Alternatively, those who settle may also be worse off too, because their outside option of going to trial is worse, forcing them to accept harsher terms.

Fortunately, other than prompting negotiations, ruling against a consumer in court is likely the only other action a judge could take that would affect financial distress. Moreover, this action is entirely observable, which means that we can simply control for it. In particular, Equation 4 and the results in Table 4 directly control for J_{ijcs}, the rate at which the judge j in court c of case i rules in favor of consumers in year s. Thus, our identification stems from the component of a judge's settlement rate that is orthogonal to the judge's consumer win rate. Panel A of Table 5 shows that our results remain are very similar whether we control for judges' in court consumer leniency.

4.3.2. What about default judgments? From Figure 1, a large proportion of cases include default judgments, or cases where the consumer fails to appear on the first appearance date despite having been successfully served. We include these in our sample to sharpen the estimates of control variables on outcomes. Nevertheless, it would be worrisome if these cases drove our results, since judge settlement propensities are supposed to reflect judge styles rather than the consumer's choice of whether to appear. To address this, Table 5 Panel B reports results where we separately control for default judgments, and the results are broadly in line with those of Table 4. This is consistent with the random allocation of judges and our identifying assumption.

4.4.3. Dynamics. Thus far we have shown that financial distress increases within the first year after settlement (specifically, by the first January after the case is disposed). We interpret this as consistent with consumers not just making bad deals that have long-term financial consequences, but bad deals that have immediate consequences. Figure 3 explores the dynamics of the effects over a longer time horizon. Specifically, we estimate Equation 4 separately in each of the 4 years before and after case disposition. As we would expect, we find no effect of in the years prior to case disposition. Interestingly, we primarily find an effect in the first year after case disposition and then point decline each subsequent year and are not statistically significant.

5. More evidence on bad deals

5.1. Prediction 2: Credit score improvement

Although the evidence above rejects Prediction 1, Prediction 2 suggests that heightened risks of delinquency, bankruptcy, and foreclosure may form part of a "good deal" if consumers benefit by avoiding damage to their credit report. Ex ante, we argue this is unlikely to be true as consumers in our sample have an average credit score of 536, which is already severely damaged.

We test Prediction 2 three ways. First, we examine whether there is a causal effect of settlement on credit scores in Equation 4. Any benefit to credit scores from settling and avoiding a judgment should show up as a positive causal effect, but Table 6 Panel A, column 1, suggests that the point estimate is +46 but statistically indistinguishable from zero. Judgments may have noisy impacts on credit scores because of difficulties with merging judgment data with credit bureau records. In 2017, all three major bureaus removed civil judgments as a factor in the computation of credit scores altogether. Notably, even a full 46-point increase does not increase the average credit score of 536 beyond 660, a widely-used industry benchmark for good credit.

Second, we examine whether there is a causal effect of settlement on mortgage and non-mortgage inquiries. The key benefit of a better credit score is that consumers can open accounts on better terms, or at all, as lenders can see this credit score through inquiries. If consumers value their credit score for this benefit, it is reasonable to expect that settlement should lead to more inquiries to a consumer's credit report compared to consumers who go to court. Table 6 Panel A, columns 2 and 3, estimate Equation 4 with the number of credit inquiries as a left-hand side variable and finds no effect of settlement on inquiries.

Finally, we show that the effects do not vary with a consumer's pre-case-resolution credit score. If the preservation or improvement of credit score was a key motive for consumers to strike deals in our sample, groups who value their credit score the most should be more willing to strike deals that risk delinquency, bankruptcy, and foreclosure. For example, high-score consumers may value their credit scores much more than lower-score consumers. Table 6 Panel B presents results where we include an additional interaction of settlement with an indicator ("Low credit score") for whether the consumer had a credit score in the lowest quartile of the distribution in the January before case disposition. We employ a similar two-stage least squares strategy where we additionally instrument the interaction of "Settlement x Low credit score" with "Judge settlement propensity x Low credit score." The interaction term is statistically insignificant in all specifications.

Overall, these results suggest that the credit scores are not a major driver of the deals we observe in our sample. Again, we emphasize that this is perhaps not too surprising given that our sample is comprised almost entirely of consumers with extremely low credit scores.

5.2. Prediction 3: Court costs as sources of surplus

Collectors who would earn a surplus from a deal (e.g., by avoiding court fees) may offer large concessions to obtain a deal if consumers have significant bargaining power. In Table 7 Panel A we examine whether consumers are able to do so.

We begin in Column 1 by again estimating Equation 4, but in this case the outcome variable is an indicator equal to one if the consumer has an account in collections. As we would expect, settling reduces the probability of having an account in collections relative to going to trial. For those who settle, the status of the settled account likely changes, as it is no longer in collections. For those who go to trial, the account likely remains in collections as the collector attempts to collect payment through garnishment. These results help to confirm that the instrument does work and that we are measuring settlement correctly.

Prediction 3 says that a good deal involving higher consumer financial distress must have settlement amounts that fall below garnishment amounts. The Missouri court data tracks garnishment cash flows, but not the dollar amount of the private settlements. Instead, our approach is to calculate garnishment cash flows and then estimate the required haircut that consumers would need to extract in a settlement to break-even with garnishment.

We start in Column 2 by examining collection balances as the left-hand side variable in Equation 4. For those who settle, the collection balance of the settled account likely goes to zero. For those who do not settle, the balance may decrease slightly to reflect the amount paid over the first year. Nevertheless, we view the difference as a proxy for the face value of debt owed by the consumer. The coefficient of $7300 is easily within one standard error of the average collection balance reported in Table 2. The sample size for this column is smaller because we limit the sample to cases where we have data on the total garnishment cash flow.

Garnishment involves an implicit haircut on average, as collectors may end up garnishing less than the full amount if the borrower declares bankruptcy, loses their job, or for other reasons. We add up all garnishment payments ever made by each defendant in connection to the case under observation and subtract this from their total collection balance the year after case disposition. We then re-estimate Equation 4 with this variable on the left-hand side. The estimated coefficient reported in Column 3 represents the dollar haircut that collectors incur from garnishment, which equals $4500, or 62% of the face value calculated in Column 2.

Consumers who settle would need to obtain a haircut of at least this amount to break even with the expected payouts from going to court. This number is conservative in that we assumed a discount rate of zero for the garnishment cash flows. To justify a good deal and compensate for the higher risk of financial distress, consumers would need to extract even larger haircuts, up to a maximum of 62% plus the collector's court cost. We view this as possible but unlikely: settlements likely involve non-trivial amounts given the subsequent financial distress, and professional collectors face arguably low marginal court costs.

5.3. Prediction 4: Risk aversion

Finally, even if making a deal leads to worse consumer outcomes on average, consumers may be risk averse, and therefore willing to accept this in exchange for a reduction of uncertainty. That is, they may be willing to bear higher expected financial distress from settlement to avoid the gamble of going to court. However, as shown earlier, most consumers lose in court; therefore, they do not actually face that much uncertainty from choosing to go to trial. Put differently, consumers are largely facing the choice of settling or losing in court and our results suggest that they are better off losing in court.

Nonetheless, in Table 7 Panel B, we explore this issue further. Since most consumers who do not settle lose in court, we simply limit the sample to only include those who settle or lose in court and then repeat our analysis. As can be seen, we continue to find similar results with this sample restriction. These results suggest even more directly that consumers are better off losing in court than settling. Therefore, even if settling offers a more certain outcome, no level of risk aversion would justify choosing it, as it offers a dominated outcome distribution.

6. External validity

How likely are our findings to apply beyond our specific setting? As stated earlier, approximately 14% of consumers with a credit file have at least one account in third-party collection each year. These consumers routinely negotiate deals with collectors, often without being litigated. While we only study those who are litigated, it seems plausible that those who are not may similarly make bad deals. In fact, one could argue that bad deals that trigger financial distress may be even more prevalent among those who are not litigated, as those who are litigated are likely selected by collectors based their ability to pay. Moreover, the types of debt collectors who operate through the court system and negotiate at the behest of a judge likely engage in fewer nefarious practices than others, leading to fewer bad deals for consumers. However, our findings suggest that consumers make bad deals even under these circumstances.

On the other hand, since we employ an instrumental variables strategy, we can only identify a local average treatment effect (LATE), or the causal effect of settlement among those for whom drawing a high settlement propensity judge is pivotal in the settlement decision. Therefore, one possible concern is that settlements may on average be good for consumers, but the settlements that judges specifically induce are bad. For that to be the case, it would have to be that consumers are actually well-informed but drawing a high settlement-propensity judge causes them to become ill-informed. Perhaps one reason that may occur is if consumers misinterpret a judge's prompt to negotiate as a signal that they are less protected under the legal system than they thought (and, in fact, are). However, if consumers' beliefs about their rights are so easily swayed, they were arguably never actually that well informed to begin with and likely could have been convinced by a debt collector to make a bad deal even without a judge's negotiation prompt.

7. Conclusion

The number of consumers experiencing debt collection dwarfs the number undergoing bankruptcy or foreclosure. Yet we know very little about the effect of debt collection on these consumers. Policymakers have long worried that many consumers may be unaware of their rights when negotiating with debt collectors and are thus vulnerable to being taken advantage of. In this paper, we investigate whether consumers who negotiate directly with debt collectors agree to deals that suggest they are ill-informed. To do so, we exploit a setting where randomly assigned civil court judges have different propensities to encourage consumers and collectors to negotiate with one another before a trial. Using new data that links court records from debt collection cases with subsequent consumer credit outcomes, we find that consumers who settle causally experience higher short-term financial distress than if their case had been resolved in court. Overall, the evidence suggests consumers likely fail to walk away from bad deals offered by debt collectors.

References

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, "Behavioral household finance." NBER working paper (2018).

Chang, Tom, and Antoinette Schoar. "Judge specific differences in Chapter 11 and firm outcomes." Working Paper (2008).

Consumer Financial Protection Bureau. "How to respond when a debt collector contacts you in three easy steps." https://www.consumerfinance.gov/about-us/blog/how-respond-when-debt-collector-contacts-you-three-easy-steps/ (2017) [Last accessed: June 2018].

Consumer Financial Protection Bureau. Fair Debt Collection Practices Act: Annual Report (2014).

Dawsey, Amanda E., Richard M. Hynes, and Lawrence M. Ausubel. "Non-judicial debt collection and the consumer's choice among repayment, bankruptcy, and informal bankruptcy." *American Bankruptcy Law Journal* 87 (2013): 1-26.

Dobbie, Will, and Jae Song. Debt relief and debtor outcomes: Measuring the effects of consumer bankruptcy protection. *American Economic Review* 105.3 (2015): 1272-1311.

Dobbie, Will, Paul Goldsmith-Pinkham, and Crystal Yang. Consumer bankruptcy and financial health. *The Review of Economics and Statistics* 99.5 (2017): 853-869.

Doyle, Joseph. "Child protection and adult crime: using investigator assignment to estimate causal effects of foster care." *Journal of Political Economy*, 116(2008): 746-770.

Doyle, Joseph. "Child protection and child outcomes: measuring the effects of foster care." *American Economic Review*, 97(2007): 1583-1610.

Dynan, Karen. “Is household debt overhang holding back consumption?” *Brookings Papers on Economic Activity* (2012): 299-362.

Ernst and Young. “The impact of third-party debt collection on the US national and state economies in 2013.” Technical report (2013).

Fedaseyeu, Viktar. “Debt collection agencies and the supply of consumer credit.” Working paper (2014).

Fedaseyeu, Viktar, and Robert M. Hunt. "The economics of debt collection: enforcement of consumer credit contracts." (2014).

Federal Reserve Bank of New York. *Quarterly report on household debt and credit*. February (2018).

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. "Liquidity constraints and consumer bankruptcy: Evidence from tax rebates." *Review of Economics and Statistics* 96.3 (2014): 431-443.

Hynes, Richard M. “Broke but not bankrupt: consumer debt collection in state courts.” *Florida Law Review* 60 (2008): 1-62.

Kling, Jeffrey. “Incarceration Length, Employment, and Earnings.” *American Economic Review*, 96(2006): 863-876.

Lee, Donghoon, and Wilbert van der Klaauw. “An introduction to the FRBY Consumer Credit Panel.” Federal Reserve Bank of New York. Technical report (2010).

Mian, Atif, and Amir Sufi. “Consumers and the economy, part II: Household debt and the weak US recovery.” Federal Reserve Bank of San Francisco Economic Letter (2011): 1-5.

Mian, Atif, Kamalesh Rao, and Amir Sufi. “Household balance sheets, consumption, and the economic slump.” *Quarterly Journal of Economics* 128.4 (2013): 1687-1726.

Mian, Atif, Amir Sufi, and Francesco Trebbi. "Foreclosures, house prices, and the real economy." *The Journal of Finance* 70.6 (2015): 2587-2634.

Silver-Greenberg, Jessica (2011). “In Debt Collecting, Location Matters.” *The Wall Street Journal*, July 18, 2011.

White, Michelle J. Economic analysis of corporate and personal bankruptcy law. No. w11536. National Bureau of Economic Research, 2005.

Zinman, Jonathan. “Household debt: facts, puzzles, theories, and policies.” *Annual Review of Economics* 7 (2015): 251-276.

Table 1. Summary statistics.

Panel A reports case-level summary statistics. "Settlement rate" is the fraction of cases per county that settled, Pr(Dismiss | Court) is the likelihood of being dismissed conditional on going to court, "Total judgment" is the amount owed to the plaintiff after judgement, "Garnishment rate" is the likelihood of a case having positive garnishment cash flows conditional on a judgment, and "Days to disposition" represents the length between filing date and disposition date. Panel B reports the characteristics of borrowers in our sample in the January before the disposition date of a case (column 1) and compares these with the overall population of borrowers (column 2) as well as the population of borrowers who declare bankruptcy (columns 3 and 4). Columns 2-4 are from Dobbie, Goldsmith-Pinkham, and Yang (2017).

Panel A. Case characteristics

	Mean	SD	Median	N
Settlement rate	0.17	0.37	0	82,218
Pr (Dismiss | Court)	0.06	0.23	0	68,516
Total judgment ($)	4,149	12,514	2,059	42,173
Garnishment rate	0.51	0.5	1	41,752
Day to disposition	88	83	59	82,218
N of cases per person	1.23	0.57	1	82,218

Panel B. Comparison with other samples

	Litigated cases	2% Random Sample of Credit Users	Bankruptcy filers	Chapter 13 bankruptcy filers
	(1)	(2)	(3)	(4)
Delinquency flag	0.756	0.148	0.413	0.675
Bankruptcy flag	0.075	0.010	0.007	0.048
Foreclosure flag	0.031	0.003	0.010	0.048
Revolving balance	7,325	6,010	13,080	10,010
Collection balance	7,012	600	1,430	2,500
Credit Score	536	740	630	580
Collection flag	0.756	0.137	0.296	0.467
Charge-off flag	0.410	0.065	0.188	0.310
Judgment flag	0.102	0.009	0.034	0.060
Lien flag	0.012	0.004	0.011	0.021
Age	42.7	48.6	43.7	44.8
Homeowner	0.510	0.470	0.520	0.643

Table 2. Borrower characteristics.

The table compares average characteristics for borrowers in the sample of cases that were heard in court (column 1) with borrowers who settled their cases (column 2). The p-value in column 3 is from a test of differences between these two samples, controlling for court by year fixed effects, where we cluster standard errors at the judge level. We compute borrower characteristics in the year before case disposition. Columns 4-6 repeat this exercise where we split cases by high or low judge settlement propensity, defined as SP_{ijcs} above or below its sample median. */**/*** denotes coefficients which are statistically reliably different from zero at the 10%, 5%, and 1% levels, respectively.

	Settlement vs Court			Judge Settlement Propensity		
	Settle	Court	p-value	Low	High	p-value
	(1)	(2)	(3)	(4)	(5)	(6)
Settlement propensity	0.006	-0.002	(0.004)***	-0.017	0.031	(0.000)***
Household distress						
Delinquency	0.55	0.52	(0.000)***	0.53	0.53	(0.400)
Bankruptcy	0.21	0.05	(0.000)***	0.08	0.07	(0.223)
Foreclosure	0.03	0.03	(0.021)**	0.03	0.03	(0.700)
Debt balances						
Revolving balance	10,164	6,758	(0.000)***	7,337	7,301	(0.534)
Mortgage balance	46,014	28,154	(0.000)***	30,722	31,946	(0.108)
Access to credit						
Credit score	562	530	(0.000)***	535	537	(0.229)
Non mtg. inquiries	1.6	1.9	(0.000)***	1.9	1.8	(0.898)
Mortgage inquiries	0.1	0.1	(0.062)*	0.1	0.1	(0.396)
Number of trade lines	4.1	3.2	(0.000)***	3.4	3.3	(0.606)
Have a collection	0.67	0.77	(0.000)***	0.76	0.76	(0.521)
Collection balance	5,776	7,259	(0.000)***	7,027	6,983	(0.397)
Have a judgement	0.05	0.11	(0.000)***	0.10	0.11	(0.546)
Have a lien		0.01	(0.005)***	0.01	0.01	(0.567)

Table 3. First-stage estimation.

The table reports the results of estimating Equation 3, where the left-hand side variable is the settlement indicator S_i and the main right-hand side variable of interest is the judge settlement propensity score, SP_{ijcs}. Column 1 reports estimates without including controls, while controls for age (binned), credit score (binned), days to disposition (binned), homeownership, pre-period bankruptcy and J_{ijcs} judge j in court c leniency are included in Column 2. Column 3 shows the first stage but controlling for age, credit score, days to disposition as a continuous variable to illustrate the impact and directions of the correlations. We cluster standard errors at the judge level and report them in parentheses. */**/*** denotes coefficients which are statistically reliably different from zero at the 10%, 5%, and 1% levels, respectively.

	Settlement indicator S_i		
	(1)	(2)	(3)
SP_{ijcs}	0.94	0.82	0.81
	(0.04)***	(0.06)***	(0.06)***
J_{ijcs}			-0.15
			(0.07)**
Credit Score			0.008
			(0.007)***
Age			-0.063
			(0.167)
Homeowner			0.046
			(0.003)***
Days to disposition			0.66
			(0.106)***
Previous bankruptcy			0.25
			(0.02)***
Controls	No	Yes	Yes
N	82,218	82,218	82,218
R^2	0.051	0.157	0.145
Clusters	43	43	43
First-stage F-statistic	737	94	160

Table 4. Second-stage estimation.
Panel A reports results from estimating Equation 4 using ordinary least squares (OLS; Columns 1, 3, and 5) and two-stage least squares (2SLS; Columns 2, 4, 6). The dependent variables in Panel A are indicator variables for delinquency, defined as a least one payment account delinquent, bankruptcy, and foreclosure, all measured in the January after case disposition. For 2SLS estimates, we instrument settlement S_i using the judge settlement propensity SP_{ijct}. Panel B repeats this exercise where the dependent variables are the dollar balance of revolving and mortgage accounts. Control variables are those defined in Section 3.1. We cluster standard errors at the judge level and report them in parentheses. */**/*** denotes coefficients which are statistically reliably different from zero at the 10%, 5%, and 1% levels, respectively.

Panel A. Financial distress

	Delinquency		Bankruptcy		Foreclosure	
	OLS	2SLS	OLS	2SLS	OLS	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)
Settlement	0.000	0.13	-0.047	0.11	-0.012	0.04
	(0.005)	(0.08)*	(0.004)***	(0.03)***	(0.001)***	(0.01)***
Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	82,218	82,218	82,218	82,218	82,218	82,218
Clusters	43	43	43	43	43	43

Panel B. Mechanism

	Revolving Balance		Mortgage Balance	
	OLS	2SLS	OLS	2SLS
	(1)	(2)	(3)	(4)
Settlement	-246	2,467	7,370	30,163
	(175)	(4,404)	(659)***	(7,141)***
Controls	Yes	Yes	Yes	Yes
N Obs.	82,218	82,218	82,218	82,218
N Clusters	43	43	43	43

Table 5. Robustness

Panel A reports results dropping the control for J_{ijcs} in Equation 4. Panel B controls for defaults judgment as an indicator. For all specifications, we estimate Equation 4 using two-stage least squares and cluster standard errors at the judge level and report them in parentheses. */**/*** denotes coefficients which are statistically reliably different from zero at the 10%, 5%, and 1% levels, respectively.

Panel A. Other Judge Actions

	Delinquency	Bankruptcy	Foreclosure
	(1)	(2)	(3)
Settlement	0.08	0.12	0.03
	(0.08)	(0.03)***	(0.01)***
Controls	Yes	Yes	Yes
N Obs	82,218	82,218	82,218
N Clusters	43	43	43

Panel B. Default Judgments

	Delinquency	Bankruptcy	Foreclosure
	(1)	(2)	(3)
Settlement	0.15	0.18	0.06
	(0.11)	(0.06)***	(0.02)***
Controls	Yes	Yes	Yes
N Obs	82,218	82,218	82,218
N Clusters	43	43	43

Table 6. Are these bad deals (1)?

Panel A estimates Equation 4 where the dependent variables are a consumer's credit score and the number of credit report inquiries (non-mortgage and mortgage) in the January after the case is disposed and where we estimate the equation using two-stage least squares. Panel B re-estimates Equation 4 for delinquency, bankruptcy, foreclosure, mortgage balances, revolving balances, and consumer credit score where we include an additional interaction of settlement with an indicator for whether the consumer had a credit score in the lowest quartile of the distribution ("Settlement x Low credit score") in the January before case disposition. We employ a similar two-stage least squares strategy where we additionally instrument the interaction of "Settlement x Low credit score" with "Judge settlement propensity x Low credit score." For all specifications, we cluster standard errors at the judge level and report them in parentheses. */**/*** denotes coefficients which are statistically reliably different from zero at the 10%, 5%, and 1% levels, respectively.

Panel A: Credit scores and inquiries

	Credit Score	Non-Mtg. Inquiries	Mortgage Inquiries
	(1)	(2)	(3)
Settlement	46	-0.55	0.01
	(33)	(0.59)	(0.02)
Controls	Yes	Yes	Yes
N Obs	82218	82218	82218
N Clusters	43	43	43

Panel B. Differential effects by credit score group

	Delinquency	Bankruptcy	Foreclosure	Credit Score
	(1)	(2)	(3)	(4)
Settlement	0.05	0.13	0.02	46
	(0.09)	(0.03)***	(0.01)**	(27)
Settlement x	0.37	-0.03	0.08	8.7
Low Credit Score	(0.25)	(0.15)	(0.07)	(29)
Controls	Yes	Yes	Yes	Yes
N Obs	82218	82218	82218	82218
N Clusters	43	43	43	43

Table 7: Are these bad deals (2)? Other surplus
Panel A estimates Equation 4 using two-stage least squares where the dependent variables are an indicator equal to 1 if the consumer has a collection flag (column 1), a consumer's collection balance in their credit report (column 2), and the consumer's collection balance minus the amount paid if they went to court and lost (column 3). For columns 2 and 3, we limit the sample to those who report a judgment and garnishment amount. We measure all credit report variables in the January after case disposition. Panel B re-estimates Equation 4 for delinquency, bankruptcy, and foreclosure, but where we exclude consumers who went to court and had their case dismissed. For all specifications, we cluster standard errors at the judge level and report them in parentheses. */**/*** denotes coefficients which are statistically reliably different from zero at the 10%, 5%, and 1% levels, respectively.

Panel A. Concession Payment Amounts

	Have a Collection	Collection Balance	Collection Bal. & Paid Amount
	(1)	(2)	(3)
Settlement	-0.35	-7,338	-4,534
	(0.09)***	(1,557)***	(1,020)***
Controls	Yes	Yes	Yes
N Obs	82218	36135	36135
N Clusters	43	43	43

Panel B. Risk Aversion

	Delinquency	Bankruptcy	Foreclosure
	(1)	(2)	(3)
Settlement	0.13	0.11	0.04
	(0.07)*	(0.03)***	(0.01)***
Controls	Yes	Yes	Yes
N Obs	78302	78302	78302
N Clusters	43	43	43

Figure 1. Litigation outcomes in Missouri
This graph plots the frequency distribution for case outcomes in the sample of counties with random judge assignment (N=82,218).

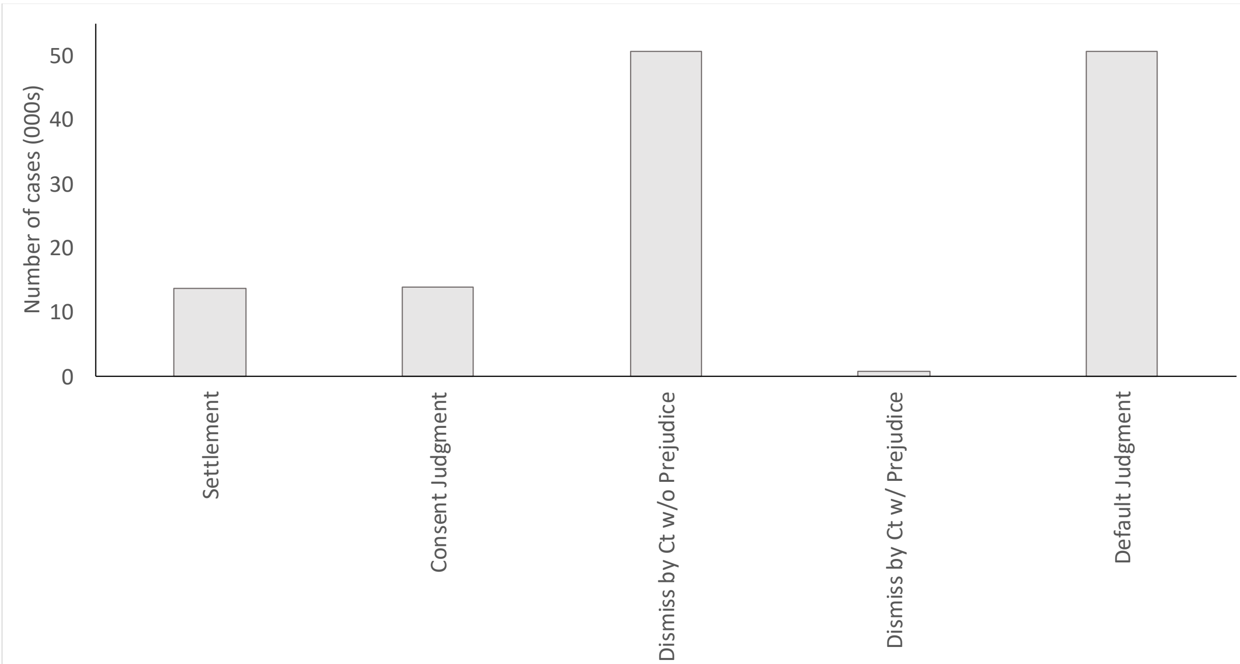

Figure 2. First Stage
This figure plots a settlement indicator vs. our leave-one-out measure of judge settlement propensity. To construct the binned scatter plot, we first regress an indicator for settlement on court-by-disposition-year fixed effects and calculate residuals. We then take the mean residual in each judge-by-year bin, adding the mean discharge rate to each residual to aid in the interpretation of the plot. The solid line shows the best linear fit estimated on the underlying microdata estimated using OLS. The coefficients show the estimated slope of the best-fit line including court-disposition year fixed effects, with standard errors clustered at the judge level reported in parentheses.

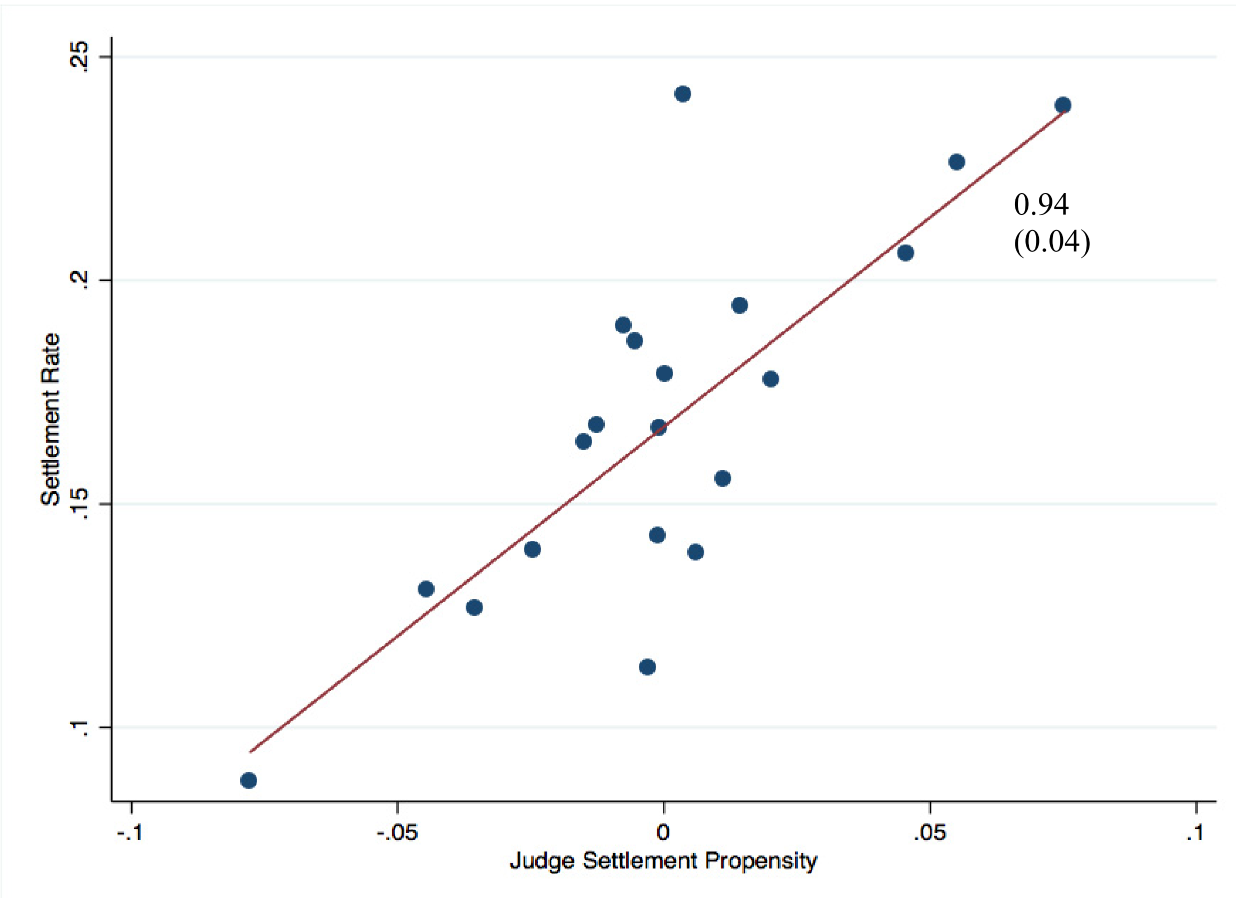

Figure 3. Dynamics

This figure reports coefficients from estimating Equation 4 for bankruptcy separately for each event year before and after case disposition. We measure event years in the January months before and after the case disposition date.

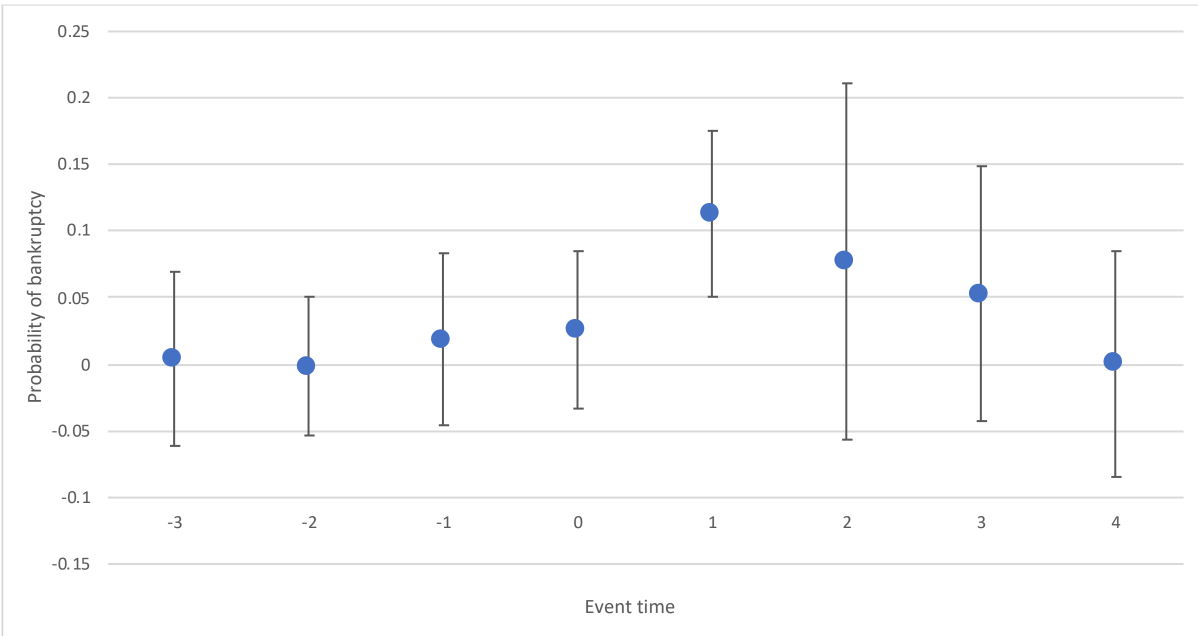

How do consumers fare when dealing with debt collectors? Evidence from out-of-court settlements

Online Appendix

Table A1. Final sample definition.
This table describes the filters we apply at each stage to arrive at our final sample.

All cases, 2007-2014	667,337
Sample of counties with random judge assignment	203,298
Matched with TransUnion in January before disposition	176,769
…match rate	87.0%
Settlement propensity measure	165,697
Settlement propensity and Matched with TransUnion	143,896
Require t=0 and t=1 presence + Data cleaning	142,038
With lawyer classification	135,989
Cases where borrower was served	82,218
Final matched sample	82,218

Table A2. Persistence of judge settlement propensity.
This table reports estimates of a regression where the dependent variable is the raw judge propensity to settle for a year, and the independent variable is the same propensity in the previous year. Column 1 reports the coefficient with robust standard errors. Column 2 reports the same coefficient but with standard errors clustered at the judge level. Standard errors are are reported in parentheses. */**/*** denotes statistically reliably different from zero at the 10%, 5%, and 1% levels, respectively

	Judge Settlement Propensity t	
Judge Settlement Propensity t-1	0.662	0.662
	(0.117)***	(0.135)***
N Obs.	141	141
R	0.297	0.297
Cluster		34

Figure A1. Geographical Distribution of Litigation in Missouri Sample
This figure shows the percent of cases that belong to each county in Missouri. The left panel shows the whole sample (N=667,337), while the right one shows the sample matched with TransUnion (N=82,218).