

2019

Manipulating Random Assignment: Evidence From Consumer Bankruptcies in the Nation's Largest Cities

Edward R. Morrison
Columbia Law School, emorri@law.columbia.edu

Belisa Pang
Yale University Law School, belisa.pang@yale.edu

Jonathon Zytnick
New York University School of Law, jzytnick@gmail.com

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship



Part of the [Bankruptcy Law Commons](#)

Recommended Citation

Edward R. Morrison, Belisa Pang & Jonathon Zytnick, *Manipulating Random Assignment: Evidence From Consumer Bankruptcies in the Nation's Largest Cities*, COLUMBIA LAW & ECONOMICS WORKING PAPER NO. 614 (2019).

Available at: https://scholarship.law.columbia.edu/faculty_scholarship/2535

This Working Paper is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact scholarshiparchive@law.columbia.edu.

Manipulating Random Assignment: Evidence from Consumer Bankruptcies in the Nation's Largest Cities*

Edward R. Morrison
COLUMBIA LAW SCHOOL

Belisa Pang
YALE LAW SCHOOL

Jonathon Zytnick
COLUMBIA UNIVERSITY
DEPARTMENT OF ECONOMICS

January 28, 2021

Abstract

Random case assignment is an important feature of court decision-making because it guards against favoritism (actual or perceived). In bankruptcy cases, it can also be important because of its effects on *ex ante* lending decisions. Despite its importance, however, random assignment is failing in several major bankruptcy courts today. We show this using data on the assignment of trustees to Chapter 7 cases in courts covering Chicago, Los Angeles, and parts of New York. These trustees audit the debtor's finances, find and liquidate assets, and police compliance with the law. We find evidence that the random-assignment process is failing in two of the courts (Chicago and New York). We dig deeper into the Chicago data and find that attorneys strategically time their case filings to avoid or attract particular trustees ("trustee shopping"). By contrast, among cases filed by debtors who have not hired attorneys ("pro se filers"), there is no case imbalance across trustees. Because they do not engage in manipulation, pro se filers—who account for the bottom decile of income and asset values among Chapter 7 debtors—are the debtors most burdened by trustee-shopping. We conclude by presenting evidence that trustee-shopping is less prevalent in Los Angeles due to critical differences in its random-assignment protocol.

JEL classification: D14, D12, G33, K35, R20

Keywords: Bankruptcy, Chapter 7, Random Assignment, Trustee Shopping

*We thank David Abrams, Thomas (Zach) Horton, Robert Jackson, Joshua Mitts, Suresh Naidu, Catherine Steege, Eugene Wedoff, and seminar participants at Columbia and the 2019 Conference on Empirical Legal Studies for their help. We are grateful for a PACER fee waiver from the Bankruptcy Courts for the Central District of California, Eastern District of New York, and Northern District of Illinois.

1 INTRODUCTION

Justice is blind. For most courts, this means that cases are randomly assigned to judges, and litigants are sanctioned when they are caught shopping for judges. In bankruptcy courts, other key actors—trustees—are randomly assigned as well. In consumer Chapter 7 cases, these trustees perform many of the core functions in a typical bankruptcy case, including holding meetings with the debtor consumer, investigating his or her financial affairs, pursuing assets in the hands of third parties, and taking other steps to maximize payoffs to creditors. Trustees are therefore assigned to cases pursuant to a “blind rotation process” administered by the Department of Justice, through its Administrative Office of the U.S. Trustee (EOUST). Details about this rotation process are kept confidential in order to deter “trustee shopping” (Executive Office of the U.S. Trustees (2012)).

In practice, justice is not blind in all bankruptcy courts. We study Chapter 7 trustee assignments in three major cities—Chicago, Los Angeles, and New York—and find evidence of trustee-shopping in two of them (Chicago and New York). We explore the phenomenon in Chicago, where the bankruptcy court granted us access to detailed case records. We find that some (but not all) attorneys monitor the flow of cases to trustees and strategically time the filing of cases to increase the probability of avoiding particular trustees who are relatively aggressive in investigating cases and pursuing creditor recoveries. These trustees are disfavored because they increase the time that an attorney must spend on a case. That time investment is unattractive because most attorneys handle large case volumes, charge flat fees, and seek to minimize the time spent per case. Trustee-shopping is sufficiently pervasive in Chicago that it changes the identity of the trustee in about ten percent of cases. The burden of this shopping is borne primarily by pro se debtors, who tend to be low-income consumers. Because they are not represented by attorneys, they typically are unaware of differences across trustees and generally do not know how to manipulate the assignment system.

Trustee-shopping presents several tradeoffs for consumers. Although trustee-shopping imposes an externality on pro se debtors, it benefits other low-income consumers who hire

attorneys. This is because more-aggressive and less-aggressive trustees differ primarily with respect to their treatment of lower-income consumers. More-aggressive trustees tend to pursue collection efforts, like keeping cases open long enough to seize debtors' tax refunds, that less-aggressive trustees forgo.¹ Thus, trustee-shopping is cost-minimizing for the attorney, but it also benefits the debtor, who avoids a closer look at his or her finances and thereby obtains a broader fresh start. Although it is unclear whether the gains to winners (who receive a broader fresh start) offset losses to losers (pro se and other filers who obtain a narrower fresh start), there is a net loss in social welfare because trustee-shopping in bankruptcy reduces creditor recoveries and raises the cost of credit for everyone *ex ante*, as we show below.

There is also a harder-to-quantify loss to the legal system, at least in the abstract. Random assignment of cases to judges has been justified by courts on the ground that it "ensure[s] greater public confidence in the integrity of the judicial process," because it "guarantees fair and equal distribution of cases to all judges, avoids public perception or appearance of favoritism in assignments, and reduces opportunities for judge-shopping."² Random assignment to Chapter 7 trustees has been justified on identical grounds.³ Our evidence shows that assignment is not random in two major cities, but we are unable to quantify the harm to public confidence.

In the sections that follow, we describe the role of trustees in Chapter 7 bankruptcy (Section 2) and present a theory of random assignment, which identifies the harms arising from trustee-shopping (Section 3). We then introduce a simple test of trustee-shopping by attorneys in three major cities—Chicago, Los Angeles, and New York (Section 4). We

¹Based on conversations with trustees, we believe that the difference in trustee behavior is driven by differences in opportunity costs. All trustees are private lawyers with their own practices, which they pursue in tandem with their trustee activities. Aggressive trustees tend to have practices focused on consumer bankruptcy; less-aggressive trustees have practices focused on business bankruptcy. Collection activities that will yield small recoveries relative to collection costs are pursued by more-aggressive trustees but overlooked by less-aggressive trustees.

²*Tripp v. Exec. Office of President*, 196 F.R.D. 201, 202 (D.D.C. 2000).

³"The United States Trustee appoints panel members to chapter 7 cases on a fair and equitable basis by utilizing a blind rotation system that includes all chapter 7 cases. A system of blind rotation avoids the appearance of favoritism and eliminates the need to make individual judgments about case assignments. Over a reasonable period of time, this system normally results in asset cases being fairly and equally distributed among the panel. Because the order of assignment is not available to the public, the 'blind' rotation also reduces the likelihood that debtors can engage in 'trustee shopping'—that is, timing the filing of a petition in order to have or avoid a specific trustee appointed to the case." Executive Office of the U.S. Trustees (2012).

turn to data from Chicago to explore the phenomenon more carefully (Section 5). Specifically, we illustrate the strategies used by attorneys to manipulate the trustee-assignment process (Section 6), quantify the number of cases affected by trustee-shopping (Section 7), and explore the causes and consequences (especially for pro se debtors) of such shopping (Section 8). We conclude by examining differences in the assignment protocols in Chicago, Los Angeles, and New York (Section 9). These differences may explain why we fail to see evidence of trustee-shopping in Los Angeles but find strong evidence of it in the other cities.

2 INSTITUTIONAL BACKGROUND

United States bankruptcy law offers two avenues of relief for financially distressed consumers. One is Chapter 7, which discharges most debts, but permits secured creditors to seize collateral (such as cars and houses) and requires the consumer to surrender other assets that are not “exempt” from collection under state or federal law. In Illinois, for example, “exempt” assets include up to \$15,000 of home equity and \$2,400 of equity in a vehicle.⁴ The other avenue is Chapter 13, which discharges a broader range of debts and allows the consumer to retain all assets but requires the consumer to complete a three- to five-year repayment plan under which all disposable income goes to creditors.

This paper focuses on Chapter 7. Broadly speaking, when a consumer (or her attorney) files a bankruptcy petition in federal district court, the case is assigned to a bankruptcy judge and a Chapter 7 trustee (hereinafter, the “trustee”). In most courts, including those studied in this paper, different assignment protocols are used for judges and trustees: Judges are assigned by the court itself; trustees are assigned by the regional United States Trustee (“UST”), a government official appointed by the United States Attorney General. There are currently 21 UST regions, covering all federal judicial districts except Alabama and North Carolina.⁵ Within each region, the UST identifies a panel of trustees and as-

⁴As in most states, these numbers double when a married couple files a “joint” bankruptcy petition.

⁵In Alabama and North Carolina, Chapter 7 trustees are assigned by Bankruptcy Administrators (“BAs”), who have been appointed by the relevant U.S. Court of Appeals. Each judicial district has its own BA. The history of the BA and UST programs is discussed in General Accounting Office (1992).

signs cases to trustees pursuant to a “blind rotation system” that, “[o]ver a reasonable period of time, . . . normally results in asset cases being fairly and equally distributed among the panel.”⁶ Trustees are typically private attorneys, not public employees. The panel of trustees for the Northern District of Illinois, for example, includes 44 private attorneys, a substantial number of whom work at large law firms (e.g., Jenner & Block) or specialize in business bankruptcies in their own practices.

The job of the trustee is to investigate the consumer’s finances, identify non-exempt assets, liquidate those assets, and distribute proceeds to creditors. A key moment in the investigation is the “341 hearing,” at which the consumer testifies under oath in response to the trustee’s questions. Although the consumer is required to submit documentation of his or her financial affairs, the 341 hearing is an opportunity to explore the completeness of the documentation and determine whether the consumer is entitled to receive property in the near future (e.g., a tax refund) or whether the consumer improperly transferred property to third parties prior to filing the bankruptcy case. The trustee has a financial incentive to identify and recover non-exempt property. He or she is paid a \$60 administrative fee for every case and, in addition, 25 percent of the first \$5,000 of non-exempt property liquidated for the benefit of creditors, 10 percent of the next \$25,000, 5 percent of the next \$950,000, and 3 percent of any amount over \$1 million. These bounties, however, are collected in a tiny percentage of cases: nationally, cases with nonexempt assets account for only about 6 percent of all cases.

Although the regional UST assigns trustees to cases using a “blind rotation system,” the system is not as “blind” as it might appear. In this paper, we focus on the Northern District of Illinois (NDIL), but our understanding is that a similar case assignment system is orchestrated nationally through the regional USTs. In the NDIL, the UST will periodically notify the court clerk that 341 hearings will be held on particular dates. For each 341 hearing date (which may span more than one day), the clerk will send out a “call,” inviting trustees to volunteer to take cases for that date. Typically, two or three trustees will be

⁶Executive Office of the U.S. Trustees (2012) at 2-4. The “blind rotation” system is subject to exceptions “on occasion.” The listed exceptions are: “1. The unique characteristics of a specific case; 2. The goal of achieving equity in the assignment of cases among panel members; 3. Suspension of a trustee from case assignments; 4. Previous service in a reopened or converted case; 5. Geographic considerations; and 6. Training for new panel members.” *Id.* at 2-4 to 2-5.

selected to be “on call.” To protect the confidentiality of the blind rotation system, neither the UST nor the court announces in advance of a meeting date the identities of the trustees on-call for that date.

Attorneys file cases using the court’s online filing system (“CM/ECF”). This system provides real-time information about case filings and trustee assignments, but it also provides a method for determining which trustees are on-call on a particular day. An attorney can log into the CM/ECF system and track recent case assignments on a given day, including the names of trustees receiving cases. If the attorney disfavors the trustees who are on-call that day, he or she can postpone a filing to a future date when more favorable trustees are on-call. Thus, while the court does not pre-announce the identities of on-call trustees for a given date, attorneys can infer the identities by watching case assignments that day.

Even if the attorney is unable to delay a case filing to another date, he or she can take steps to minimize the probability that the case is assigned to a particular on-call trustee. As we show below, the UST appears to use a form of “batch assignment,” assigning sequences of filings—sometimes ten cases in a row—to each on-call trustee. At the same time, the CM/ECF system permits attorneys to file batches of cases simultaneously. This is made possible by a two-step procedure for filing cases: First the attorney uploads cases, each of which is immediately assigned a case number (“docket number”) by the CM/ECF system. The cases are not assigned to a judge or trustee, however, until the attorney directs the system to assign the cases. Thus, an attorney can monitor the case assignments in real time and submit a batch of cases when the attorney believes that a favorable trustee will be receiving a batch. Additionally, because the UST’s assignment process tends to balance the caseloads of trustees, an attorney can predict whether a particular trustee will receive cases based on the relative caseloads.

3 A THEORY OF RANDOM ASSIGNMENT

Although random assignment of judges and trustees is generally viewed as a highly desirable feature of our legal system, simple economic intuition suggests that it is rarely

optimal to allocate a scarce resource randomly instead of assigning it to the individuals who value it the most. However, in the context of assignment of Chapter 7 trustees, this intuition would lead to those having the most assets to hide matching with the least aggressive trustees, which would raise interest rates and reduce welfare due to adverse selection into bankruptcy. Instead, the model suggests that those with greater willingness to pay for less aggressive trustees should not be more likely receive one. The model also shows that the mechanism for assigning trustees should not be predictably correlated with any aspect of the debtor known at the time of borrowing, so as to avoid adverse selection into borrowing. It becomes clear that random assignment of trustees (or a system closely resembling it) is optimal under conditions that characterize the Chapter 7 process today, in which the court has little information about the debtor's likelihood of having discoverable assets (since all information is self-reported).

To demonstrate these points, we use a simple two-period model with one debtor and a competitive market of creditors.

3.1 *The Debtor*

We assume a debtor with low wealth in period 1 (which we normalize to zero) and uncertain wealth in period 2. Specifically, period 2 wealth will equal y_l with probability $1 - p$ and y_h with probability p , where $0 < y_l < y_h$.

The debtor's ex ante utility is given by: $U(x_1, x_2) = u(x_1) + Eu(x_2)$ where u is a concave period utility function and x is the debtor's consumption. Because the debtor has low wealth in period 1 but higher wealth in period 2, she benefits from borrowing against future income and smoothing consumption over time.

3.2 *The Creditors*

Assume that creditors are competitive, risk-neutral, and do not discount future cash flows. In period 1, creditors offer the debtor a menu of options, each of which consist of a loan of size b and associated interest rate $r = r(b)$. The debtor selects her optimal b , and

in period 2, either repays the loan or defaults, in which case creditors receive all of the debtor's income.

We make the following assumption about the optimal loan (Assumption 1): $y_l < b < br \leq y_h$. This implies that the debtor can only pay her income when it is low (y_l), but can repay the loan (br) when her income is high (y_h).

3.3 Benchmark: Perfect Debt Enforcement

In this section, we assume that there is Perfect Debt Enforcement (PDE), with non-monetary utility cost of default c . For any given loan size b , the interest rate r_{PDE} is determined by the break-even condition for a competitive market of creditors:

$$pbr_{PDE} + (1 - p)y_l = b \quad (1)$$

In a competitive market, the debtor will choose loan size b to maximize consumer welfare:

$$U(b) = u(b) + pu(y_h - br_{PDE}) + (1 - p)(u(y_l) - c) \quad (2)$$

Because u is concave, there is a unique solution to this problem, b_{PDE} , which implies a corresponding interest rate $r_{PDE} = r_{PDE}(b_{PDE})$. By Assumption 1, $b_{PDE} > y_l$, so that utility with a loan is greater than utility without. Note that if $c = 0$, for any loan satisfying Assumption 1, even one that is not the optimal size b_{PDE} , the borrower is better off with a loan.⁷ With perfect debt enforcement, the debtor's *ex ante* utility from a loan of size b is given by:

$$U_{PDE}(b) = u(b) + pu(y_h - br_{PDE}(b)) + (1 - p)(u(0) - c) \quad (4)$$

⁷Utility with a loan is greater than utility without a loan if and only if:

$$u(b) + pu(y_h - br) + (1 - p)u(0) > u(0) + pu(y_h) + (1 - p)u(y_l). \quad (3)$$

To see that the left hand side is greater, denote β_b the slope of $u(\cdot)$ from 0 to b ; β_{y_l} the slope of $u(\cdot)$ from 0 to y_l ; and β_{br} the slope of $u(\cdot)$ from $y_h - br$ to y_h . Due to the concavity of $u(\cdot)$ and the definition of $r = r(b)$, it follows that $\beta_{br} < \beta_b < \beta_{y_l}$. This implies that the difference between consumer utility with a loan and without a loan is equal to:

$$\begin{aligned} [u(b) - u(0)] - p[u(y_h) - u(y_h - br)] - (1 - p)[u(y_l) - u(0)] \\ = \beta_b b - p\beta_{br}br - (1 - p)\beta_{y_l}y_l \\ = b(\beta_b - \beta_{br}) + (1 - p)y_l(\beta_{y_l} - \beta_{br}) > 0. \end{aligned}$$

3.4 *Imperfect Debt Enforcement*

Henceforth we assume the debtor has some power over the repayment process. Specifically, we assume that, at the start of period 2, the debtor obtains private information about her wealth in period 2 (y_h or y_l). If it is low, she defaults and files for bankruptcy. If it is high, she has a choice: She can either repay the loan or default and file for bankruptcy.

The bankruptcy process has several important features. First, a debtor choosing to enter the bankruptcy process takes on a non-monetary utility cost equal to c . Second, after filing, the debtor pays the creditors whatever assets she claims to own, but there is no penalty for lying about what she owns (or, equivalently, it is prohibitively costly for the trustee to prove perjury). Thus, a high-income debtor will always claim to have low income y_l . Third, the bankruptcy court may order an audit of the debtor's financial affairs in order to find assets that have not been disclosed.

Because it is infeasible to audit every debtor, the portion of debtors audited is fixed at α , where $0 \leq \alpha \leq 1$. We interpret α as the proportion of trustees who are "aggressive" in the sense that they have a positive probability of discovering assets during an audit, if assets exist. The remaining trustees $(1 - \alpha)$ are "lax" trustees who never discover assets, even if they exist.

3.5 *Random Audits*

With this structure in place, we can now assess how the mechanism by which audits are assigned affects welfare. Suppose audits are probabilistic, with all debtors having an equal likelihood uncorrelated with any known aspect. Given audit probability α , a high-income debtor will repay her loan (and not file for bankruptcy) if:

$$\alpha \geq \underline{\alpha} \equiv 1 - \frac{c}{u(y_h - y_l) - u(y_h - br(b))}.$$

We assume that $\alpha \geq \underline{\alpha}$: that is, the audit probability is high enough to deter cheating. If $\alpha < \underline{\alpha}$, then the audits would be ineffective and the courts would be strictly better off having no audits at all.

If all debtors are randomly assigned audits with probability $\alpha \geq \underline{\alpha}$, then debtors who realize a high income will repay their loans without entering bankruptcy, and those who realize a low income will enter bankruptcy at utility cost c and pay y_l . The break-even condition for the interest rate $r_R(b)$ will be given by $(1 - p)(y_l) + pbr_R = b$, which is the same interest rate as the perfect debt enforcement rate $r_{PDE}(b)$. Similarly, optimal loan size $b_R = b_{PED}$ and the debtor's utility (and thus social welfare) is the same as in PDE:

$$U_R(b_R) = u(b_R) + pu(y_h - b_R r_R(b_R)) + (1 - p)(u(0) - c) = U_{PED}(b_{PED}) \quad (5)$$

3.6 Auctioned Audits

What if a debtor can affect her own probability of being audited in bankruptcy? Conditional on entering bankruptcy, high-income debtors will be willing to pay any price below $y_h - y_l$ to avoid audit, and low-income debtors receive no benefit and will not pay to avoid audit at any price.

Because there are only two levels of income in our simple model, the auction will not fully determine whether a given debtor will be audited (unless the proportion of low income debtors $1 - p$ precisely equals the proportion of audits, α). Suppose that there are at least as many low-income debtors as there are audits, $1 - p \geq \alpha$. Then all high-income debtors will enter bankruptcy and purchase non-audits at some positive price, and the probability of a low-income debtor being audited will be $\frac{\alpha}{1-p}$.

In this scenario, the credit market breaks down: Creditors can never receive back more than y_l , and so they never lend more than y_l . Social welfare is reduced:

$$U_A(b = y_l) = u(y_l) + pu(y_h - y_l) + (1 - p)u(0) < U_R(b_R) \quad (6)$$

If all high-income debtors are able to avoid an audit, credit markets will collapse beyond the minimum level of income. If only a fraction of high-income debtors are able to avoid an audit, interest rates will increase, but markets may continue to function. In either

circumstance, social welfare is lower than in a world where audits are assigned randomly.⁸ The situation can be described as adverse selection into trustees: The debtors willing to pay for trustees are precisely those for whom it is most socially valuable to audit. High-income debtors will pay to avoid audits; only low-income debts will be subject to them, and these audits will rarely discover assets.

Conversely, debtors are *ex ante* better if debtors who are *less* willing to pay for less aggressive trustees are assigned less aggressive trustees. The results of this subsection imply that, to the extent information about debtors is known prior to trustee assignment, that information should be used to assign trustees. But if little is known about debtors *ex ante*—as is typically the case in Chapter 7 cases—a debtor’s willingness to pay for a trustee should not influence the selection of trustees.

3.7 Audits Assigned by Some Other Factor

The previous subsection does not prove that it is optimal for all debtors to have an equal likelihood of receiving an aggressive trustee. We have shown only that a system in which debtors pay for trustees is suboptimal relative to one with simple random assignment. By “simple random assignment” we mean assignment based on a characteristic (e.g., case number) that is observable at the time of the bankruptcy but is (a) uncorrelated with the debtor’s wealth and (b) unobservable at the time of the lending decision. In this subsection, we modify our model to describe the characteristics of an optimal assignment mechanism. We do that by comparing simple random assignment to any other method of random assignment that is based on a characteristic known at the time of the lending decision.

We begin by varying the information environment in period 1 by assuming that, when a debtor applies for a loan, she is characterized by a type that is (a) unrelated to her income in period 2 but (b) related to her probability of being audited in period 2. Specifically, assume a fraction p_θ of debtors have type θ ; the remaining debtors have type $\bar{\theta}$. This type

⁸Note that paying to avoid being audited is therefore socially inefficient even if no effort is expended to do the actual manipulation. To the extent the attorney must expend resources to do the manipulation, that adds to the inefficiency.

is unrelated to the debtor's probability of being high- or low-income in period 2, but it determines the probability of an audit in bankruptcy: All debtors with type $\bar{\theta}$ have zero probability of audit; all debtors with type θ have probability of audit equal to $\frac{\alpha}{(1-p\theta)}$.

We consider two cases. Suppose, first, that the debtor's type is *public* information in period 1. In this case, there will be a breakdown of the credit market for those with type $\bar{\theta}$, reducing social welfare as compared to the random assignment case. Next suppose that the debtor's type is *private* information (unobservable to the creditor). Now we are in an adverse selection model, and creditors will be forced to offer smaller loans with higher interest rates than in the random assignment model.

This shows that tethering audit assignment to anything known at the time of borrowing, public or private, induces adverse selection into borrowing and reduces social welfare. While we have not proven that random assignment is optimal, we have shown that optimal assignment must be uncorrelated with any factors known at the time of borrowing and must not be determined by willingness to pay for an audit.

The key takeaways from our model are as follows: 1) Social welfare is greater when higher-income debtors have a higher probability of audit. 2) Higher-income debtors are willing to pay more for less aggressive trustees. Thus, a system in which trustees are awarded by auction is welfare-reducing. 3) If the court does not know the debtor's income at the time of trustee assignment, as in the current Chapter 7 process, then any mechanism which is correlated with something known to the debtor at the time of borrowing is welfare-reducing.

4 TRUSTEE-SHOPPING: A SIMPLE TEST

The discussion in Section 2 suggests that trustee-shopping can occur along two margins: Attorneys can avoid filing cases on days when disfavored trustees are on-call ("extensive margin") and, conditional on filing cases during a particular day, attorneys can time their filings to reduce the probability of assignments to disfavored trustees ("intensive margin"). Evaluating shopping along the intensive margin requires data on the exact time of

day when each case was filed; such data are difficult to obtain, for reasons discussed below. Evaluating shopping along the extensive margin requires data on the days when cases were filed; these data are relatively easy to obtain because each bankruptcy court produces summary case-level data, indicating the days on which cases were filed and the identities of the attorneys, judges, and trustees associated with each case.⁹ We obtained these data for the years 2008 through 2017 from three courts—Central District of California (“CDCA,” which includes Los Angeles), Northern District of Illinois (“NDIL,” which includes Chicago), and Eastern District of New York (“EDNY,” which includes parts of New York City and all of Long Island).¹⁰ Using these data, we can (i) identify, for each trustee, the days when the trustee was on-call and (ii) calculate, for each attorney, the percentage of on-call days when the attorney did not file cases. This calculation is a potential measure of trustee-shopping (attorneys avoiding certain trustees). However, attorneys may not file cases when a trustee is on-call for reasons independent of trustee-shopping, such as caseload, fluctuations in caseload, personal vacations, etc. For example, attorneys that file a small number of cases per year inevitably will not file cases on many days when certain trustees are on-call.

We address these issues by focusing on days within weeks during which the attorney and trustee were both active. Specifically, we identify the top twenty law firms in each court, as measured by annual caseload, and calculate, for each firm-trustee pairing, the percentage of *weeks* when (a) the trustee was on-call at least once, (b) the firm filed at least one case during the week, and (c) the firm never filed on a day when the trustee was on-call. We plot the distribution of this measure of trustee-shopping and compare it to a simulated distribution, which calculates the percentage of weeks when an attorney would not match to a trustee simply due to random variation in trustee assignments and the flow of cases. We create the simulated distribution as follows: For every week during which a firm filed cases, we “refile” those cases on a randomly selected day of the week. If a trustee was not on-call that day, we say the firm “avoided” the trustee on that day. Of course, this avoidance is purely due to random chance. We then calculate, for each attorney-trustee

⁹These data are called the “Cases Report” and are available for purchase from each court, or can be accessed without charge if the court grants a fee waiver.

¹⁰We use data from the Los Angeles Division of the CDCA, Brooklyn and Long Island Divisions of the EDNY, and Eastern Division of the NDIL.

pair, the percentage of weeks when the trustee was on-call, the firm filed at least one case, and yet none of the firm's cases were assigned to that trustee. We regenerate the simulated distribution 1,000 times for each court.

Figure 1 presents this measure of trustee-shopping. Comparing the actual and simulated distributions, we see a sharp divergence in NDIL (Chicago) and EDNY (New York), with a substantially higher probability that an attorney will avoid a trustee than is predicted by random variation. By contrast, we see no evidence of trustee-shopping in CDCA (Los Angeles), where the actual distribution is virtually identical to the simulated.¹¹

To be sure, the patterns in Figure 1 could be attributable to the UST assignment protocol, not to attorney behavior. The UST admits that it will make exceptions to the "blind rotation" system "on occasion."¹² Reasons for making exceptions include the "unique characteristics of a specific case" and to achieve "equity in the assignment of cases among panel members."¹³ The UST "documents the reasons for an exception to the blind rotation system." *Id.* Based on discussions with practitioners and court officials, we believe that all cases in Chicago are subject to the blind rotation system except extremely large cases (roughly speaking, with assets of \$100 million or more) and re-filings of previous cases. In the sections that follow, we explore the extent to which attorney behavior, not simply the UST protocol, explains the non-randomness observed in Figure 1. We do this using data for Chicago (NDIL), where we were able to obtain more comprehensive data than for other locations.¹⁴ In the conclusion, we discuss reasons why trustee-shopping is evident in Chicago and New York but not Los Angeles.

5 DATA FOR CHICAGO

We study trustee-shopping in Chicago by linking four data sources. One is the Integrated Database (IDB) maintained by the Federal Judicial Center.¹⁵ The IDB provides

¹¹These distributions are constructed based on pairings of law firms and trustees. We obtain similar results when we construct the distributions based on lawyer-trustee pairings.

¹²Executive Office of the U.S. Trustees (2012).

¹³*Id.* at 2-4 to 2-5.

¹⁴We were unable to obtain sufficiently broad fee-waivers from the other courts in order to conduct comparable analysis.

¹⁵The IDB is available online at www.fjc.gov/research/idb.

case-level summaries for all bankruptcy cases closed or filed between October 2007 and October 2017.¹⁶ The summary information includes consumer assets, debt, and monthly income; case duration; and case outcome (whether the consumer's debts were discharged or the case was dismissed). Our second data source is the "Cases Report" available from the NDIL's CM/ECF site. As noted above, this report includes the name and address of the consumer, the name and address of the consumer's attorney (if an attorney was hired), and the names of the judge and trustee assigned to each case. Our third data source is information scraped from the NDIL's CM/ECF site. For each case, we obtained the exact time when docket numbers were assigned to cases and when the judge and trustee were assigned to each case. Additionally, we downloaded the bankruptcy petition, financial schedules, and final trustee report. Together, these documents allow us to identify the assets that each consumer reported at filing and the types of nonexempt property available for liquidation. Finally, we impute each debtor's race and ethnicity, based on his or her name and address, using census data and the database described in Tzioumis (2018). Details on the imputation procedure are available in Morrison, Pang, and Uettwiller (2020).

Table 1 summarizes our linked database, which includes about 180,000 Chapter 7 cases by consumers with an average annual income of about \$32,000. Only 1 percent of cases had nonexempt assets, lower than the national average; put differently, nothing was paid to creditors—and no bounty was paid to the trustee—in 99 percent of cases (we refer to cases in which assets were distributed as "distribution cases"). Nine percent of consumers filed for bankruptcy without an attorney. Among cases with attorneys, the caseload per law firm was about 25 cases per year, though the standard deviation is very large. One firm, in particular, accounts for 13.5 percent of all filings; the top three firms account for nearly 30 percent. Caseload per trustee averaged about 700 cases per year.

¹⁶Data from earlier years is not as readily available; it must be acquired through an application to the Inter-university Consortium for Political and Social Research (ICPSR) and is subject to confidentiality controls that make it more difficult to analyze.

6 TRUSTEE-SHOPPING IN CHICAGO: ATTORNEY STRATEGIES

6.1 *Evidence of Non-Random Assignment*

A now-large literature exploits random judge assignment in bankruptcy cases as an instrument to study the effects of judge background on case outcomes and the effects of case outcomes on post-bankruptcy consumer welfare. Examples include Dobbie and Song (2015), Dobbie, Goldsmith-Pinkham, and Yang (2017), Bernstein, Colonnelli, and Iverson (2019), Bernstein et al. (2018), and Iverson et al. (2020). As this literature discusses, random assignment should produce “balance” across judges in average case characteristics. Table 2 implements one measure of balance for the Chapter 7 cases filed in NDIL. We regress case characteristics, such as whether the consumer is African American, on a set of judge (or trustee) fixed effects plus month-year dummies. Each row of Table 2 reports output from a separate regression, predicting the case characteristic in the first column. Columns 1 and 2 report *F*-tests (and significance levels) for the joint significance of the judge (or trustee) fixed effects. Although *F*-tests are biased in favor of rejecting the null, as Abrams, Bertrand, and Mullainathan (2013) explain, Table 2 shows a clear pattern: We generally cannot reject the null hypothesis of random judge assignment, but there is strong evidence of imbalance across trustees. For example, there is no statistically significant difference across judges in the value of property, but there is a highly significant difference across trustees. Evidence consistent with shopping by attorneys is presented in Columns 3 and 4, which limits the sample to cases filed by pro se debtors. These are debtors who have not hired lawyers; they are laypeople who are representing themselves. In this subsample, where attorneys are absent, we see balance across both judges and trustees. There is no statistically significant difference across judges or trustee in the (log) value of property held by pro se debtors.¹⁷

Table 2 points to strong evidence of trustee-shopping by attorneys. We now explore the mechanisms by which this is achieved. We begin by presenting illustrative examples

¹⁷We obtain the same results using the method advocated by Abrams, Bertrand, and Mullainathan (2013) (see Table A.1 for detailed results): We can reject the null hypothesis of random trustee assignment among cases generally, but we cannot reject the null among pro se cases.

of attorney manipulation and then present a more formal measure, which allows us to say more precisely the extent to which manipulation affects case outcomes.

6.2 *Attorney Strategies*

We start with the two largest law firms, which we label as “Attorney 6” and “Attorney 17”.¹⁸ For each firm, we calculate the share of filings assigned to each Chapter 7 trustee by year. Figure 2 shows this calculation, focusing on the three trustees with the largest shares (“favored” trustees) and the three with the smallest shares (“disfavored” trustees) of each firm’s cases. Among Attorney 6’s cases, we observe large differences in trustee shares, especially after 2009. In 2013, for example, favored trustees each received about 6 percent of Attorney 6 filings, while disfavored trustees each received about 1 percent of filings. This difference persists across time. By contrast, among Attorney 17 cases, the difference between favored and disfavored trustees is much smaller and not persistent. In some years, for example, favored trustees received smaller shares than disfavored trustees. Evidently, some law firms engage in trustee-shopping more aggressively than other firms, and this shopping can lead to very large differences in case shares by trustee.

The bottom panels of Figure 2 focus on the trustees who were most-favored (Trustee 28) and least-favored (Trustee 18) by Attorney 6. Within each trustee’s cases, we rank law firms by their shares of the trustee’s cases and identify the top ten firms (thus, the “top ten” is different for each trustee). We then plot the shares of the top-ten firms as well as the shares of pro se debtors (“Pro Se”) and all other law firms (“Others”). As can be seen in Panels (c) and (d), Attorney 6 is a top-ten attorney even among cases assigned to its least-favored trustee (Trustee 18) because the firm files such a high volume of cases that it cannot avoid Trustee 18 completely. Yet Attorney 6’s share of Trustee 18 cases is much smaller than its share of Trustee 28 cases. The firm’s trustee-shopping imposes an externality on other bankruptcy filers, especially pro se consumers, who account for over 20 percent of Trustee 18’s cases in 2014, but only about 10 percent of Trustee 28 cases during that year. This points to a distributional consequence of trustee-shopping: It has greatest

¹⁸We anonymize the law firms and trustees in our discussion here, but the identities are available upon request.

impact on bankruptcy filers who are not represented by attorneys and therefore are likely unaware of differences across trustees.

So how does a firm like Attorney 6 manipulate the system? There are at least two strategies: (i) filing cases on days when a disfavored trustee is not on-call (extensive margin), and (ii) filing on days when a disfavored trustee is on-call, but timing the case filing to avoid that trustee (intensive margin).

6.2.1 Trustee-Shopping Along the Extensive Margin

We already presented evidence of “extensive margin” shopping in Figure 1, which showed how often attorneys are able to evade a trustee completely during a week when both the attorney and the trustee are actively filing and receiving cases. We see extensive-margin shopping in Chicago and New York, but not Los Angeles. However, Figure 1 aggregated the behavior of many law firms. We take a closer look at attorney behavior in Figure 3, which shows how particular firms implement extensive-margin shopping. The top panels show the number of cases filed per day by a law firm (Attorney 21) during 2013 and 2015, when we observe substantial trustee-shopping by this firm. Red dots indicate days when a particular trustee (Trustee 20) was on-call. The bottom panels show the same information for a different law firm (Attorney 22). The firms had comparable case loads per year, but there is much stronger correlation between Attorney 21’s filing behavior and Trustee 20’s on-call dates (0.17) than between Attorney 22’s behavior and Trustee 20’s on-call dates (0.06). Put differently, Attorney 21 filed cases on 66 percent of the days when Trustee 20 was on call. The comparable percentage for Attorney 22 is 41 percent.

6.2.2 Trustee-Shopping Along the Intensive Margin

Figure 4 illustrates the “intensive margin” strategy—if a firm has chosen to file a case on a day when a disfavored trustee is on-call, the firm can time the filing to reduce the probability of being assigned to that trustee. Recall from Section 2 that the UST assignment process has two features that render it vulnerable to manipulation: (i) Cases are assigned in batches and (ii) attorneys can monitor the assignment process in real time, allowing

them to see which trustees are on-call and how many cases they have received that day. Figure 4 shows how attorneys take advantage of these features. In each panel, we follow the sequence of 60 cases filed before and after the first case assigned to a group of disfavored (or favored) trustees. We focus on the two largest law firms: Attorney 6 (panels a and b) and Attorney 17 (panels c and d). We say that a trustee is disfavored if she is one of the three trustees with the smallest shares of Attorney 6 (or Attorney 17) cases, as in Figure 3. We define favored trustees similarly (i.e., those with the largest shares of an attorney's case filings). In each panel, Case 0 identifies the first case assigned to disfavored (or favored) trustees on a particular day. We then identify the 30 cases that were assigned to trustees (*any* trustees) before Case 0. We also identify the 30 cases assigned to trustees (*any* trustees) after Case 0. For example, Case -20 identifies the twentieth case assigned to a trustee *before* the first case was assigned to a disfavored/favored trustee. Similarly, Case 10 is the tenth case *after* the first case assignment to a disfavored/favored trustee.

Each figure has two lines. The red line shows the probability (measured on the right axis) that a case in this sequence is assigned to disfavored (or favored) trustees. By construction, this probability is (a) equal to 0 for cases that were assigned before a disfavored (or disfavored) trustees receive their first assignment and (b) equal to 100 percent for the first case (Case 0) assigned to those trustees. The red line illustrates the “batching” in the UST assignment algorithm. After a disfavored (or favored) trustee receives a case (Case 0), the probability of receiving the next case in the sequence is over 50 percent. The probability remains over 20 percent until about the eighth case, at which point the probability flat-lines at about 20 percent. This pattern characterizes all trustee assignments, regardless of whether the trustee is disfavored or favored (compare Panels (a) and (b)), indicating that “batching” is an artifact of the assignment process.

The blue line in each panel of Figure 4 shows the probability, by sequence number, that the case was filed by Attorney 6 (or Attorney 17). In Panel (a), for example, the probability that Attorney 6 filed Case -10 is about 15 percent, which is unsurprising since Attorney 6 accounts for about 15 percent of all cases in the sample. Notice, however, that Attorney 6's probability falls sharply—from 15 percent to about 7 percent—immediately when a disfavored favored trustee is assigned his or her first case (Case 0). The probability

does not return to its previous level (15 percent) until about Case 10, when a disfavored trustee is likely no longer receiving a “batch” of cases. When we look at favored trustees in Panel (b), we also see a decline at Case 0, but the decline reverses itself immediately afterward. This suggests that Attorney 6 reduces filings whenever any trustee is receiving a first assignment and then either reduces filings further if the trustee is revealed to be a disfavored type or resumes filing if the trustee is revealed to be a favored type.

Panels (c) and (d) perform the same exercise for the Attorney 17 firm. Recall from Figure 2 that Attorney 17 appears to engage in much less trustee-shopping than Attorney 6. This is confirmed by Panels (c) and (d): In contrast to Attorney 6, we see no meaningful change in Attorney 17’s filing behavior after disfavored trustees receive their first assignments. But we do see a change in behavior after a favored trustee receives an assignment, with Attorney 17’s share of cases rising from about 4 percent to 6 percent, and then declining to 4 percent again by Case 10. This suggests that, although Attorney 17 is a relatively less aggressive trustee-shopper, the firm may nonetheless be exploiting the batch-assignment system.¹⁹

7 QUANTIFYING MANIPULATION

Thus far we have documented the existence of trustee-shopping, but we have not quantified its magnitude. We do that here by deriving a formal test of the extent to which observed trustee assignments deviate from what would be observed under random assignment.

Assume there are A attorneys indexed by a , and J trustees indexed by j . Let N be the total number of cases filed in a court during a given period, N_a the number of cases filed by attorney a , and N_j the number of cases assigned to trustee j . Finally, let N_{aj} measure the number of attorney a cases assigned to trustee j . If cases are assigned randomly, the expected value of trustee j ’s share of attorney a ’s cases, $\frac{N_{aj}}{N_a}$, should equal trustee j ’s share of all cases, $\frac{N_j}{N}$, during the same period. Equivalently, the expected number of attorney a

¹⁹Because we condition on *ex post* outcomes—namely, we use the attorney’s most- and least-favored trustees—we would expect a small jump at case 0 even if the attorney were not engaging in any intentional trustee shopping.

cases assigned to trustee j during a given period, N_{aj} , should be equal to $N_a \cdot \frac{N_j}{N}$, and the difference should tend towards zero in large samples:

$$D_{aj} = N_{aj} - E[N_{aj}] = N_{aj} - N_a \cdot \frac{N_j}{N} \quad (7)$$

D_{aj} measures the extent to which a trustee's observed share of an attorney's cases deviates from what would be expected under random assignment. We call this a "difference score." Although the score is equal to zero in expectation under the null hypothesis of random case assignment, the observed value \hat{D}_{aj} could differ from zero due to random variation. We can therefore compare the observed value, \hat{D}_{aj} , to the values of D_{aj} that emerge from simulations that hold fixed the caseloads of the attorney, trustee, and court— N_a , N_j , and N —but randomly assign each attorney's cases to trustees that either were (a) on-call during the five days preceding the date when the attorney filed the cases or (b) were on-call on the same day when the attorney filed the cases. Simulation (a) assumes that attorneys may delay filing a case by up to five days in order to avoid a particular trustee. Simulation (b) assumes that attorneys cannot delay the filing day, but can strategically time a case filing down to the minute.

Figure 5 presents the distribution of \hat{D}_{aj} as well as the distribution of simulated D_{aj} under simulations (a) and (b). We generate 100 simulations and present the mean as well as the 5th and 95th percentiles. Both Figures 5(a) and (b) show that the distribution of \hat{D}_{aj} has fatter tails and is shifted left relative to the simulated distributions. Integrating under the curves in Figure 5(a), we find that the number of cases under the observed distribution (\hat{D}_{aj}) is 6,231 greater than the number of cases under the mean simulated distribution. Put differently, if we assume that attorneys can delay filing cases by up to five days, trustee-shopping caused 6,231 cases to be assigned to trustees that would not have received those cases under random assignment. This represents 10 percent of all cases. Using the 5th and 95th percentiles of the simulated distributions, we can say that between 5,921 and 6,524 cases (9.5 to 10.5 percent of cases) were affected by the trustee-shopping. If we assume that attorneys cannot delay filing cases, as Figure 5(b) does, trustee-shopping caused between 4,751 and 5,299 (7.6 to 8.5 percent of cases) to be assigned to trustees who would not

otherwise have received those cases.

Trustee-shopping, therefore, had a substantial effect on the allocation of cases to trustee, affecting 7.6 to 10.5 percent of cases.

8 CAUSES AND CONSEQUENCES OF NON-RANDOM ASSIGNMENT

In this section, we explore reasons why attorneys engage in trustee-shopping. We focus on differences across trustees that appear to induce attorneys to avoid or seek out particular trustees. When attorneys do this, they impose an externality on debtors whose attorneys do not engage in trustee-shopping and on debtors who proceed without attorney representation. These debtors are more likely to be assigned to trustees who are “disfavored” by the attorneys who engage in trustee-shopping. Although the externality is redistributive, simply shifting cases between trustees, it functions like a regressive tax: Higher-income debtors are represented by trustee-shopping attorneys, who avoid the trustees most likely to conduct an aggressive audit of the debtor’s finances and find assets that can be distributed to creditors, including unpaid tax refunds. These trustees are, instead, substantially more likely to be assigned cases filed by pro se debtors, who are among the poorest debtors and who generally do not engage in trustee-shopping. We document this externality at the end of this section.

8.1 *Causes*

Trustees differ along several dimensions that may induce “trustee shopping” by attorneys. Of course, observable differences across trustees are affected by the manipulation itself. We can, however, measure un-manipulated trustee differences using cases filed by pro se debtors. There is no evidence that pro se debtors engage in manipulation. Indeed, Table 2 shows much stronger evidence of balance across trustees among pro se cases than among cases in the full sample. Figure 6 limits the sample to pro se cases and shows substantial variation across trustees along several dimensions, including case duration (ranging from about 100 days to over 150); how often the trustee adjourns a Section 341 meeting, which

delays the case (from ten to thirty percent of cases); the percentage of cases with at least two motions filed by the trustee (from about five percent to over twenty), and the percentage of cases with distributions to creditors (from less than one-quarter of a percent to over six percent). Figure 7 shows that these differences persist in the full sample as well.

These differences are correlated with the extent to which attorneys engage in trustee-shopping to avoid trustees, as Table 3 illustrates. Here, we rank attorneys (and trustees) by the extent to which they are engaging in (or are the subject of) trustee-shopping. To do this, we leverage the “difference score,” D_{aj} , defined above. Recall that D_{aj} measures the difference between (i) the observed number of attorney a cases assigned to trustee j and (ii) the expected number, assuming random assignment. Although difference scores can be positive or negative, the sign of the score does not indicate the extent of trustee shopping with respect to a particular trustee. For example, if an attorney aggressively avoids a particular “disfavored” trustee, the difference score for that attorney-trustee pair will be highly negative, but the difference scores for pairings of that attorney with other trustees will be positive. This is because the attorney is steering cases to those trustees in order to avoid the disfavored trustee. Similarly, if an attorney aggressively pursues a particular “favored” trustee, the difference score with that trustee will be highly positive, but the scores with other trustees will be negative. Thus, if we want to rank attorneys by the extent to which they engage in trustee-shopping, we cannot use the sign of the difference scores; nor can we use the mean of those scores because they will cancel out. Instead, we use the absolute values of the difference scores and, for each attorney, calculate the mean of these absolute values: $\sum_j \frac{|D_{aj}|}{N_a}$. We call this the attorney’s “mean (absolute value) difference score”; it is analogous to a standard deviation because it measures the average “spread” between the attorney’s observed and expected caseload per trustee. Attorneys who engage in no trustee-shopping will have low average spreads; aggressive shoppers will have high average spreads.

Our goal is to compare the characteristics of trustees who are disfavored to those who are favored. To do this, we select the three attorneys with the highest mean (absolute value) difference scores. These are the most aggressive trustee-shoppers. Among cases assigned to these attorneys, we identify the three trustees with the most positive and negative mean

difference scores. Here we do *not* average the absolute values of the difference scores. Instead, we average the *raw* values because we want to identify the trustees who these attorneys avoid, on average.

Table 3 shows that, among aggressive trustee-shoppers (Panel A), the disfavored trustees have relatively high case durations, are substantially more likely to file three or more motions per case and are much more likely to find assets that can be distributed to creditors, especially tax refunds. We view this as evidence that trustee-shopping is driven by attorneys seeking to avoid trustees who are, on average, more likely to conduct longer investigations of the debtor's assets. These investigations burden the attorney by consuming time that could be allocated to other cases. Our own informal conversations with bankruptcy attorneys and trustees confirm that attorneys disfavor trustees who conduct lengthy investigations.

8.2 Consequences

The consequences of trustee-shopping are clear when we look at pro se cases. Because they are non-strategic, pro se filers are more likely to be assigned to cases that the attorneys actively avoid. This is clear in the first row of Table 3, which shows that pro se cases account for 13.8 percent of the caseload of disfavored trustees, but 9.7 percent of other trustees' caseload. Trustee-shopping by attorneys, in other words, imposes an externality on pro se filers, pushing them toward disfavored trustees. Panel C shows that, among pro se cases, debtors assigned to disfavored trustees are more likely to be forced to distribute assets to creditors, and much more likely to have their tax refunds seized for distribution. Because pro se debtors have substantially lower income and assets than other debtors, as Table 1 shows, trustee-shopping effectively results in a regressive tax: The lowest income (and lowest wealth) debtors are the most likely to be forced to make payments to creditors.

9 DISCUSSION AND CONCLUSION

We find evidence of non-random case assignment in two major bankruptcy jurisdictions covering Chicago and parts of New York City. Using detailed data from Chicago, we

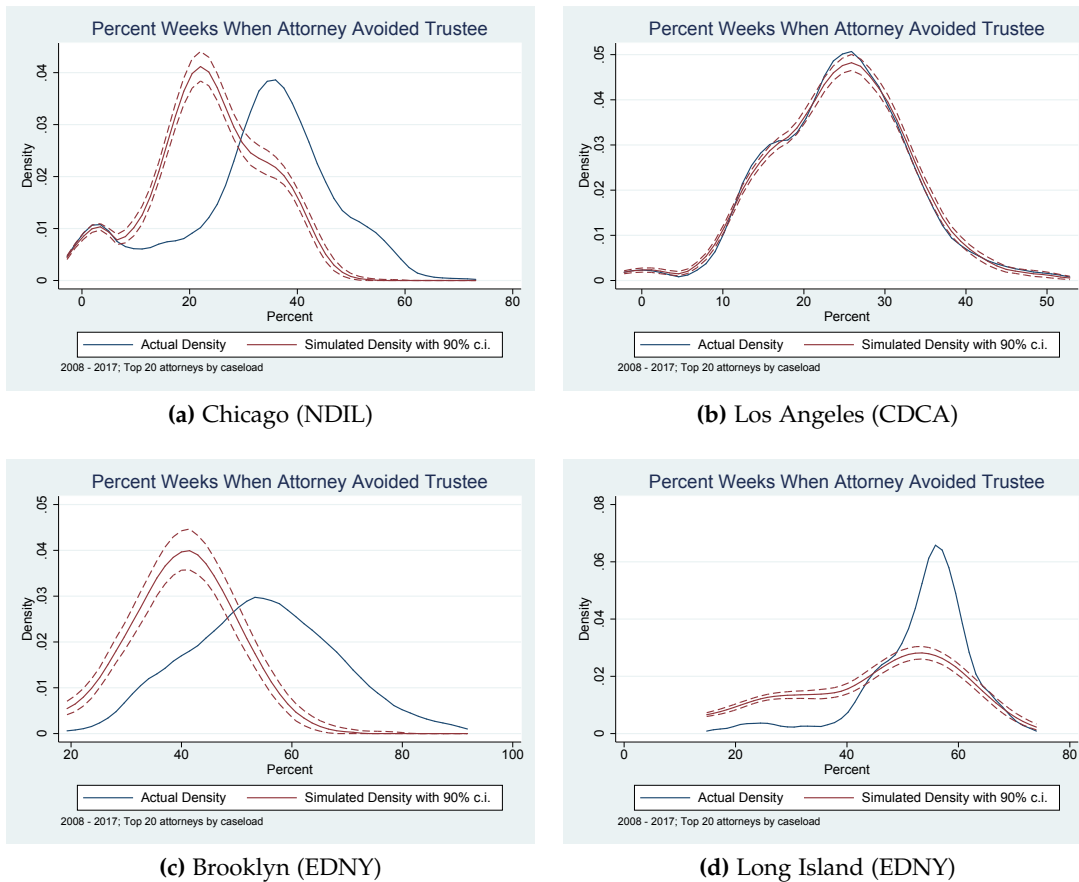
show that the non-randomness is produced by attorneys engaged in trustee-shopping. The attorneys exploit predictable patterns in the court's assignment process, such as putting a small number of trustees (two or three) on call at a time and assigning cases in batches (averaging eight cases per batch). Shopping affects about 10 percent of all cases filed in Chicago and imposes an externality on pro se debtors, who are indirectly steered toward trustees who are most aggressive in finding and liquidating assets for the benefit of creditors.

A puzzle in our data is the absence of trustee-shopping in Los Angeles. This puzzle finds a partial explanation when we compare the case assignment protocol in L.A. versus the other courts. Although all trustees in all courts are assigned using the "blind rotation system" of the regional USTs, the protocol works differently across courts, as shown by Figure 8 and Table 4. Relative to the other courts, L.A. is less likely to assign cases in batches (the median batch size is one) and its batches are smaller (L.A.'s mean batch size is about 2; for the other courts, it is between 6 and 10). The smaller batch sizes make it more difficult to shop along the "intensive margin," i.e., wait to file a case until a favored trustee has begun to receive a batch of cases (see Section 6.2.2 above). Additionally, L.A. trustees are more likely to serve multiple on-call dates in a row: The average trustee serves on about 4 consecutive on-call dates; in other courts, the average trustee serves on 2 consecutive dates. When trustees serve on multiple on-call dates, it's more difficult for attorneys to shop on the "extensive margin," i.e., wait for a day when a disfavored trustee is not on-call (see Section 6.2.1 above). Additionally, on-call dates span multiple days (about 3) in L.A., but last only one day in other courts ("Batch Length (days)"). This too reduces attorneys' ability to shop along the extensive margin because the same trustee will be present multiple days in a row. In future work, we plan to test whether such inter-court differences explain can explain the apparent absence of trustee-shopping in L.A.

REFERENCES

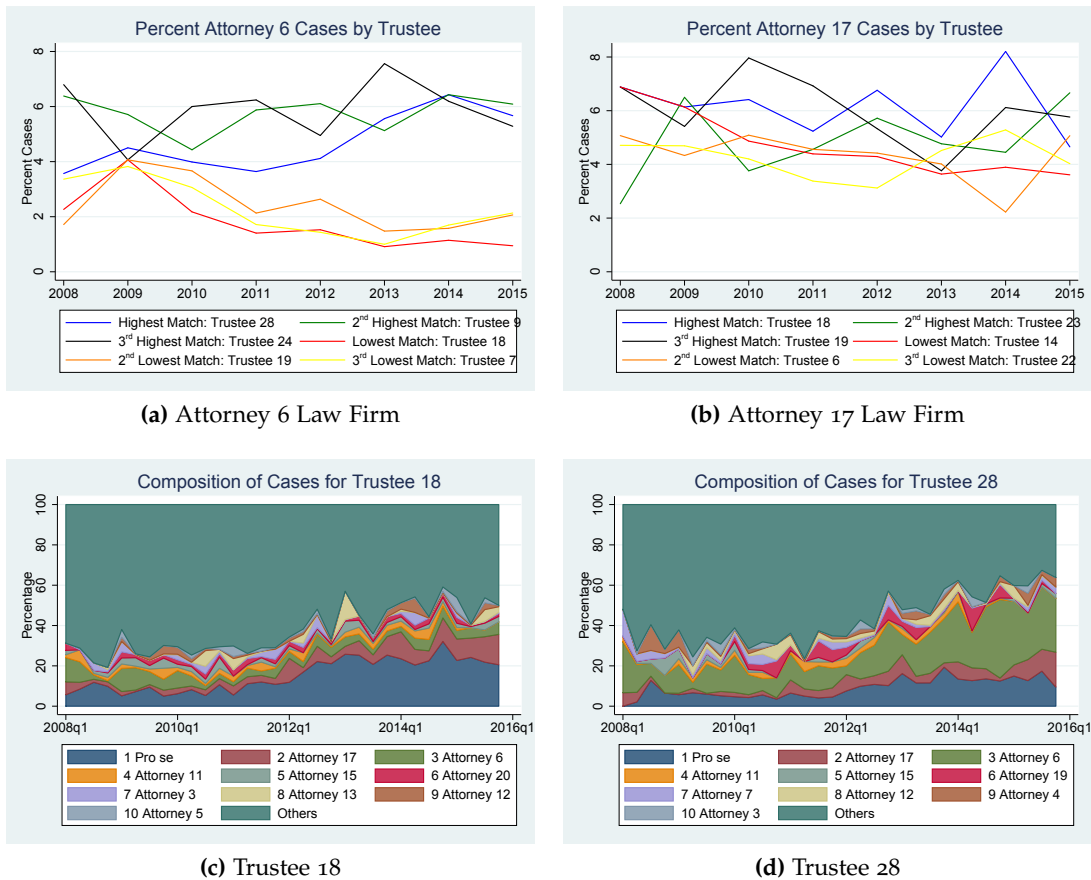
- Abrams, David S., Marianne Bertrand, and Sendhil Mullainathan. 2013. "Do Judges Vary in Their Treatment of Race?" *Journal of Legal Studies* 41:347–83.
- Bernstein, Shai, Emanuele Colonnelli, Xavier Giroud, and Benjamin Iverson. 2018. "Asset Allocation in Bankruptcy." *Journal of Finance* 74 (1):5–53.
- Bernstein, Shai, Emanuele Colonnelli, and Benjamin Iverson. 2019. "Bankruptcy Spillovers." *Journal of Financial Economics* 133 (3):608–633.
- Dobbie, Will, Paul Goldsmith-Pinkham, and Crystal Yang. 2017. "Consumer Bankruptcy and Financial Health." *Review of Economics and Statistics* 99 (5):853–69.
- Dobbie, Will and Jae Song. 2015. "Debt Relief and Debtor Outcomes: Measuring the Effects of Consumer Bankruptcy Protection." *American Economic Review* 105 (3):1272–1311.
- Executive Office of the U.S. Trustees. 2012. "Handbook for Chapter 7 Trustees." URL https://www.justice.gov/ust/file/handbook_for_chapter_7_trustees.pdf/download.
- General Accounting Office. 1992. "Bankruptcy Administration: Justification Lacking for Continuing Two Parallel Programs." No. B-248877.1, September 28, 1992.
- Iverson, Benjamin, Joshua Madsen, Wei Wang, and Qiping Xu. 2020. "Financial Costs of Judicial Inexperience." Working paper, BYU Marriott School of Management.
- Morrison, Edward R., Belisa Pang, and Antoine Uettwiller. 2020. "Race and Bankruptcy: Explaining Racial Disparities in Consumer Bankruptcy." *Journal of Law and Economics* 63 (2):269–95.
- Tzioumis, Konstantinos. 2018. "Demographic aspects of first names." *Scientific Data* 5 (www.nature.com/articles/sdata201825).

FIGURE 1: Evidence of Trustee-Shopping, by Court



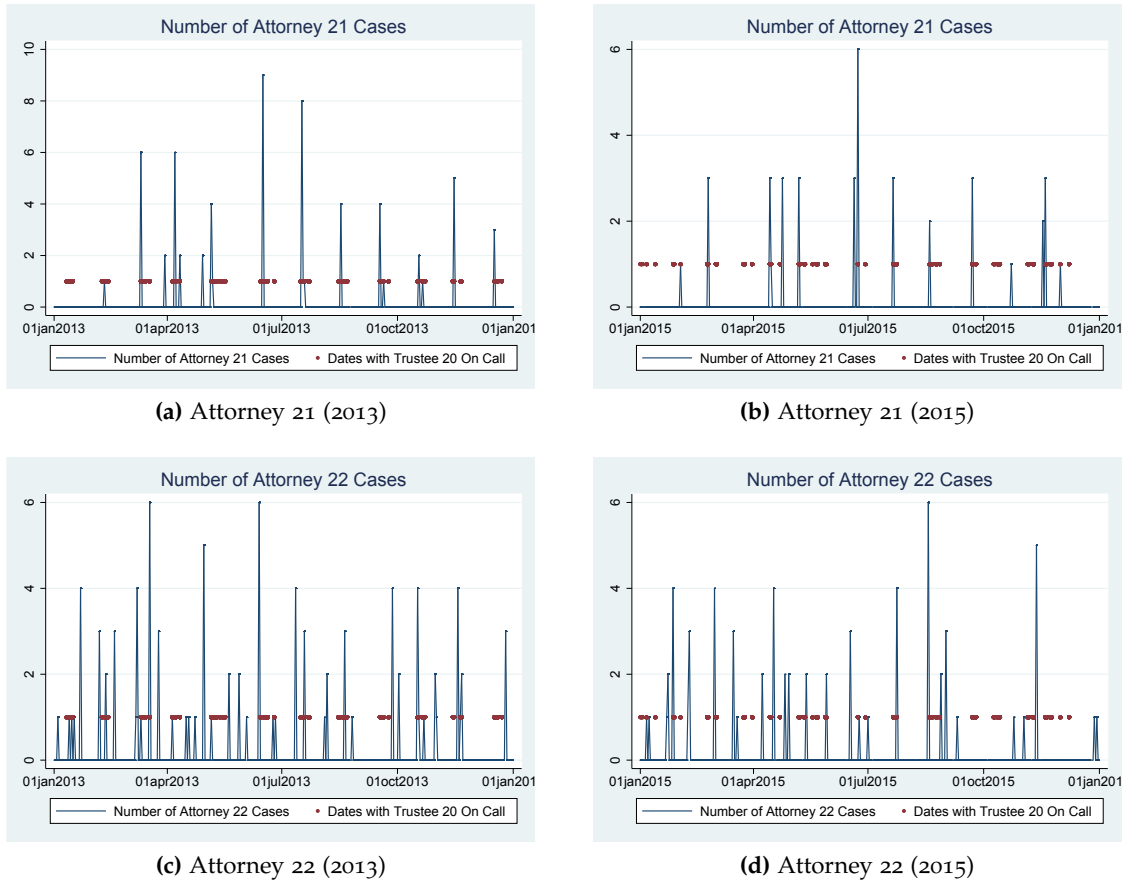
This figure contains, for four courts, a measure of the extent of attorney-trustee non-matching. Specifically, for the top twenty law firms in each court by annual caseload, we calculate, for each firm-trustee pairing, the percentage of weeks when (a) the trustee was on-call at least once, (b) the firm filed at least one case during the week, and (c) the firm never filed on a day when the trustee was on-call. The blue line contains the actual distribution of this measure of trustee-shopping and the red line contains a simulated distribution. We create the simulated distribution as follows: For every week during which a firm filed cases, we “refile” those cases on a randomly selected day of the week. If a trustee was not on-call that day, the firm “avoided” the trustee on that day. We then calculate, for each attorney-trustee pair, the percentage of weeks when the trustee was on-call, the firm filed at least one case, and yet none of the firm’s cases were assigned to that trustee. We regenerate the simulated distribution 1,000 times for each court.

FIGURE 2: Trustee Share of Cases, by Attorney by Year



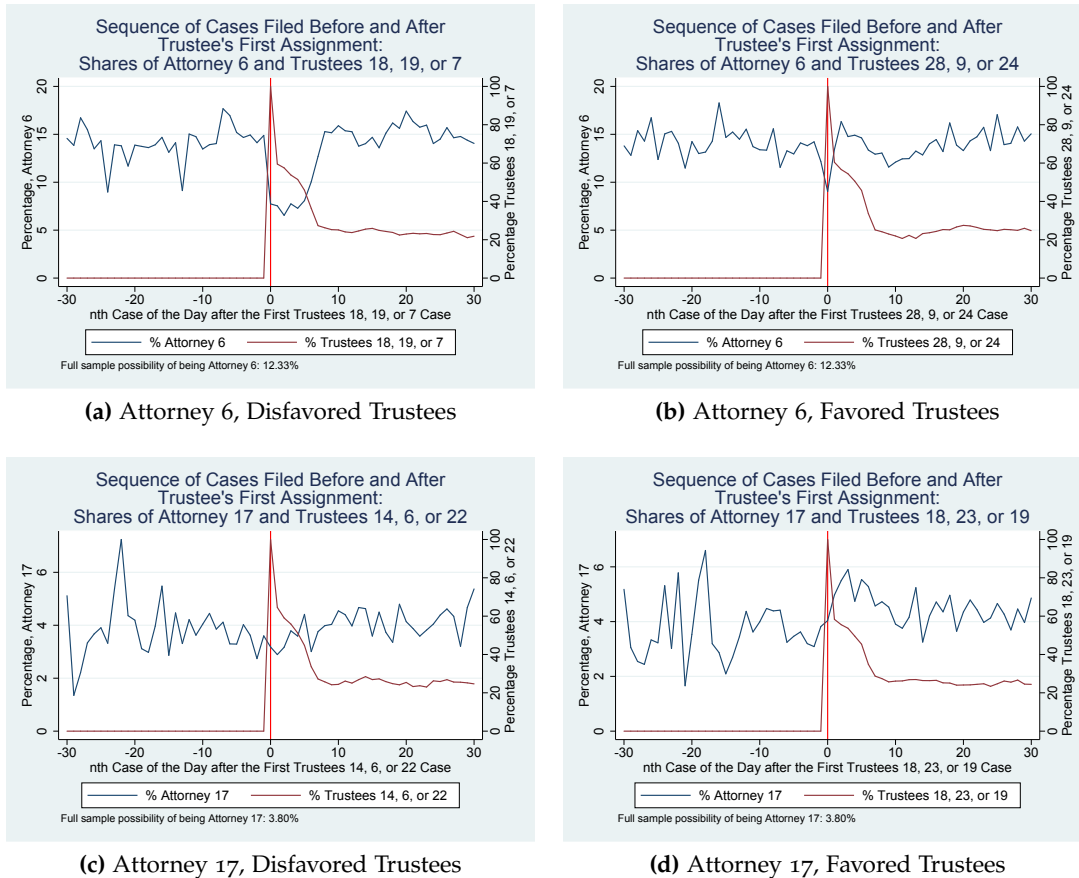
This Figure presents shares of cases in which certain law firms match with certain trustees. The top panels contain, for the two largest law firms (labeled “Attorney 6” and “Attorney 17”), the share of filings assigned to certain Chapter 7 trustees by year. These panels are limited to the three trustees with the largest shares and the three with the smallest shares of each firm’s cases. The bottom panels contain, for the trustees who were most-favored (Trustee 28) and least-favored (Trustee 18) by Attorney 6, the share of filings from certain law firms each year. Within each trustee’s cases, we rank law firms by their shares of the trustee’s cases and identify the top ten firms (thus, the “top ten” is different for each trustee). We then plot the shares of the top-ten firms as well as the shares of pro se debtors (“Pro Se”) and all other law firms (“Others”).

FIGURE 3: Attorney Filings, by Trustee On-Call Dates



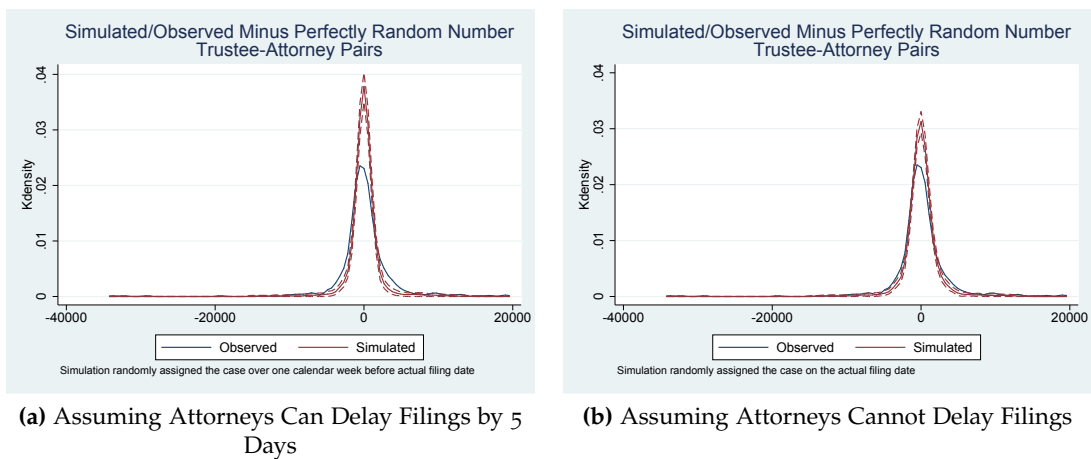
Each plot of this Figure contains, by date, the number of filings by a certain law firm and an indication of whether Trustee 20 was on call on that date.

FIGURE 4: Attorney’s Probability of Filing a Case in the Sequence Preceding and Following the First Case Assigned to Favored and Disfavored Trustees



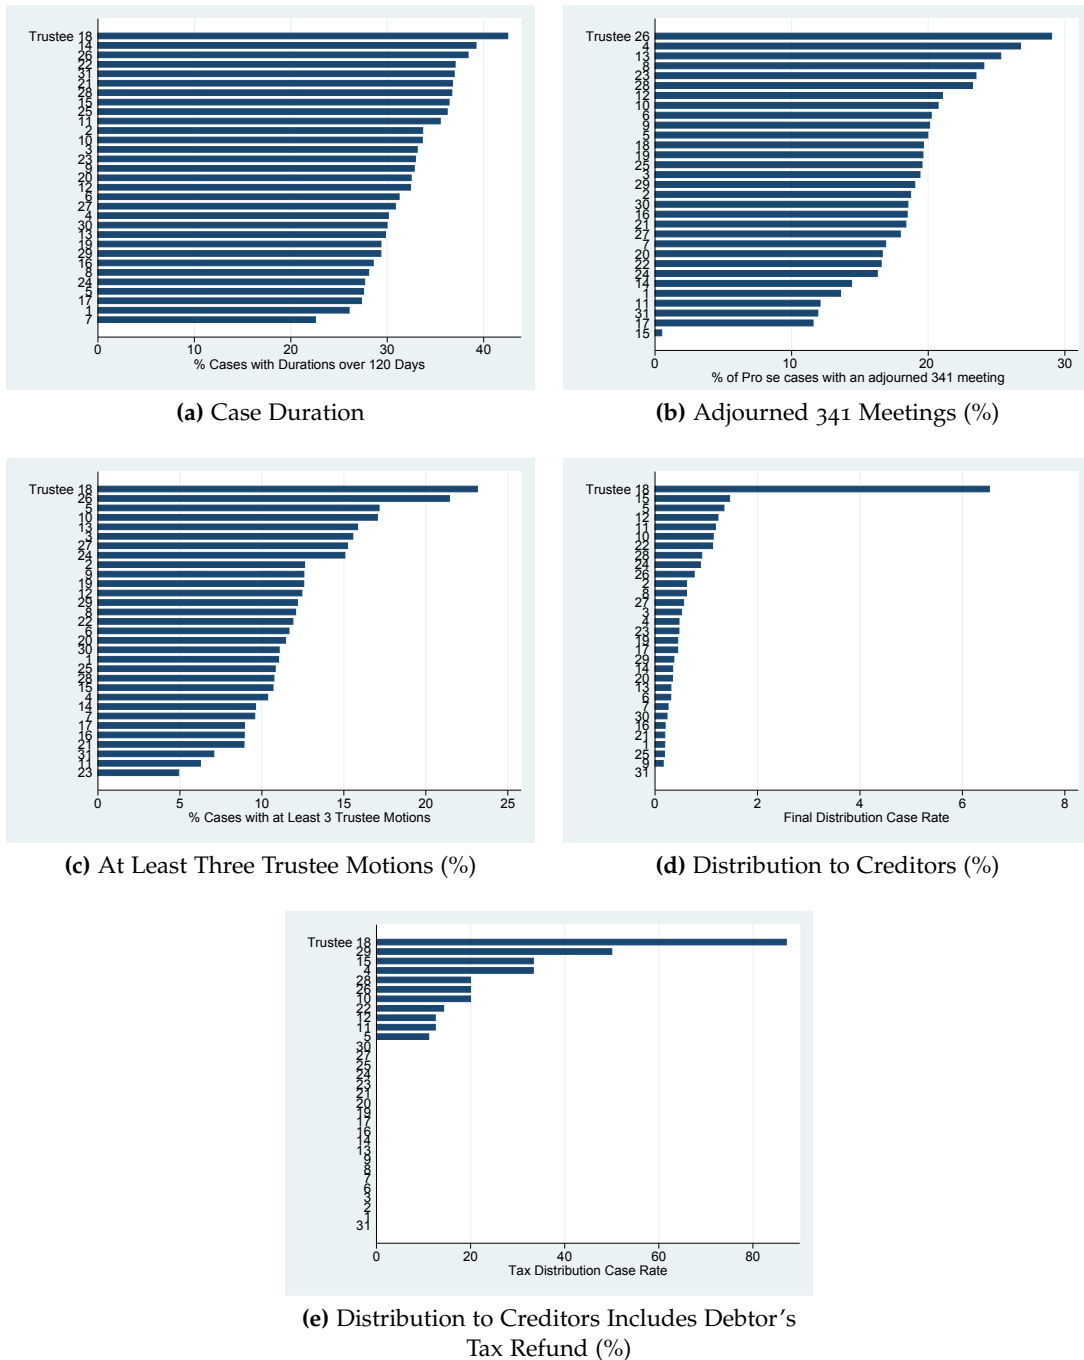
This Figure contains within-day probabilities of law firm case filings and trustee case assignments. The blue line in each plot represents the probability that a given case is assigned to a certain law firm; the red line in each plot represents the probability that a given case is assigned to any of three trustees. Each plot displays the sequence of 60 cases filed before and after the first case assigned to one of the three trustees included in the plot. Case 0 on the x-axis identifies the first case assigned to the group of trustees on a particular day. The top plots display results for Attorney 6; the bottom plots display results for Attorney 17. The left (right) plots display results for the three trustees comprising the smallest (largest) shares of the law firm’s cases.

FIGURE 5: Difference Between Observed and Expected Trustee-Firm Caseloads:
Actual v. Simulated



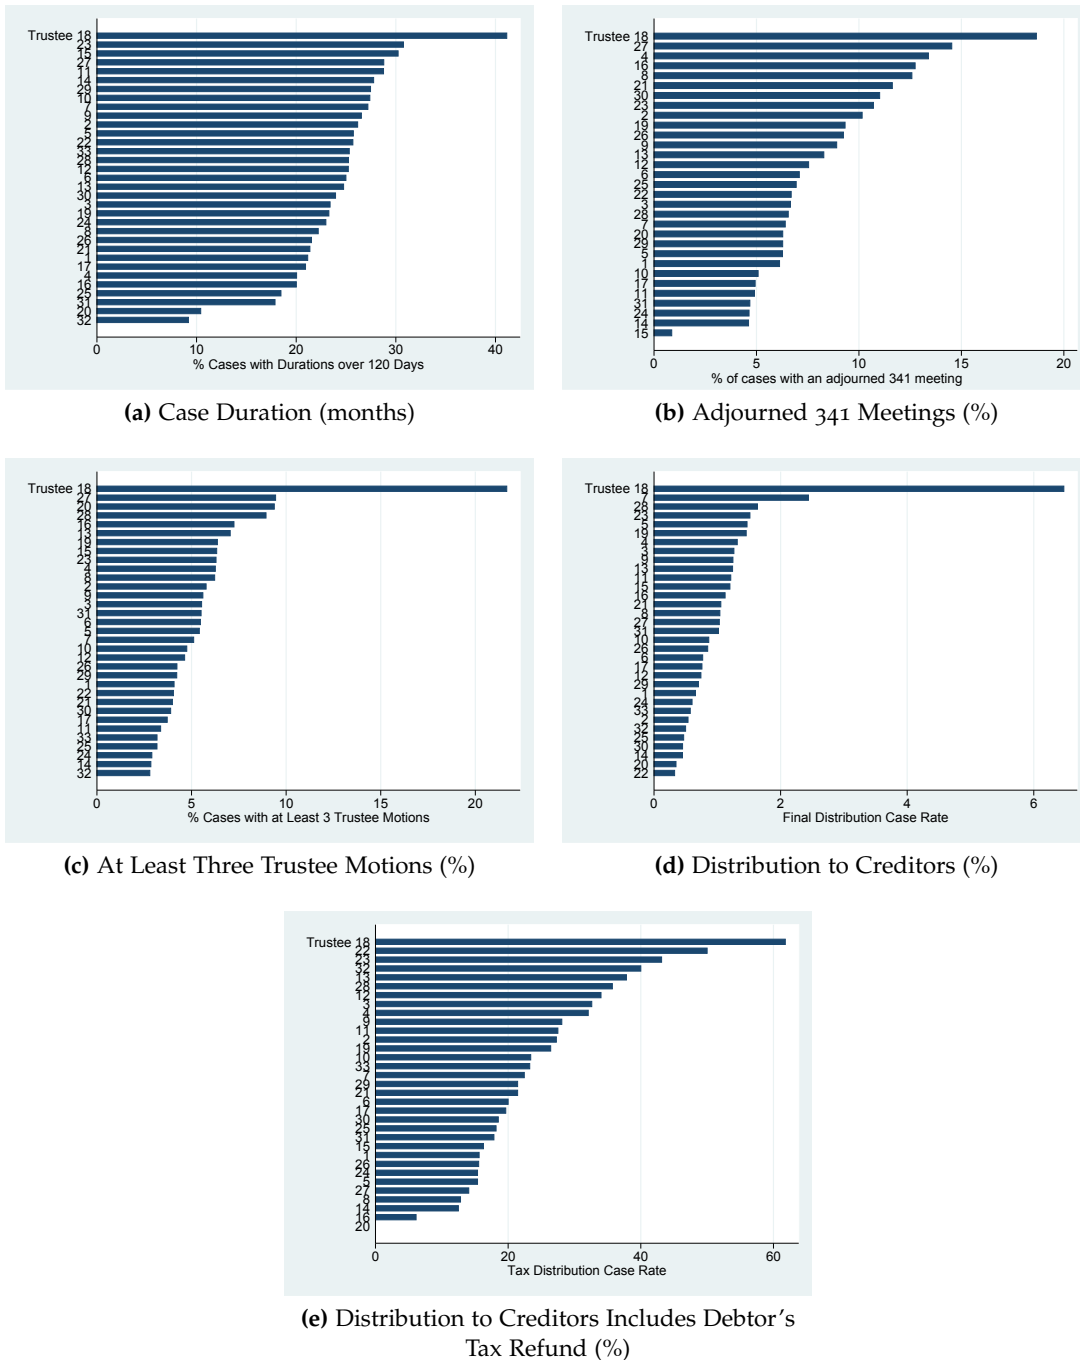
Each panel of this Figure presents the kernel density distribution of D_{aj} , the difference between the number of matches between attorney a and trustee j and the expected number of matches, defined as the product of attorney a 's number of cases and trustee j 's number of cases divided by the total number of cases. The blue lines represent the actual distribution, and the red lines represent simulated distributions with a 90% confidence interval. The simulation in Panel (a) randomly re-assigns cases within a 5-day period, and the simulation in panel (b) randomly re-assigns cases within a single day.

FIGURE 6: Variation in Decision-Making Across Trustees, Pro Se Cases



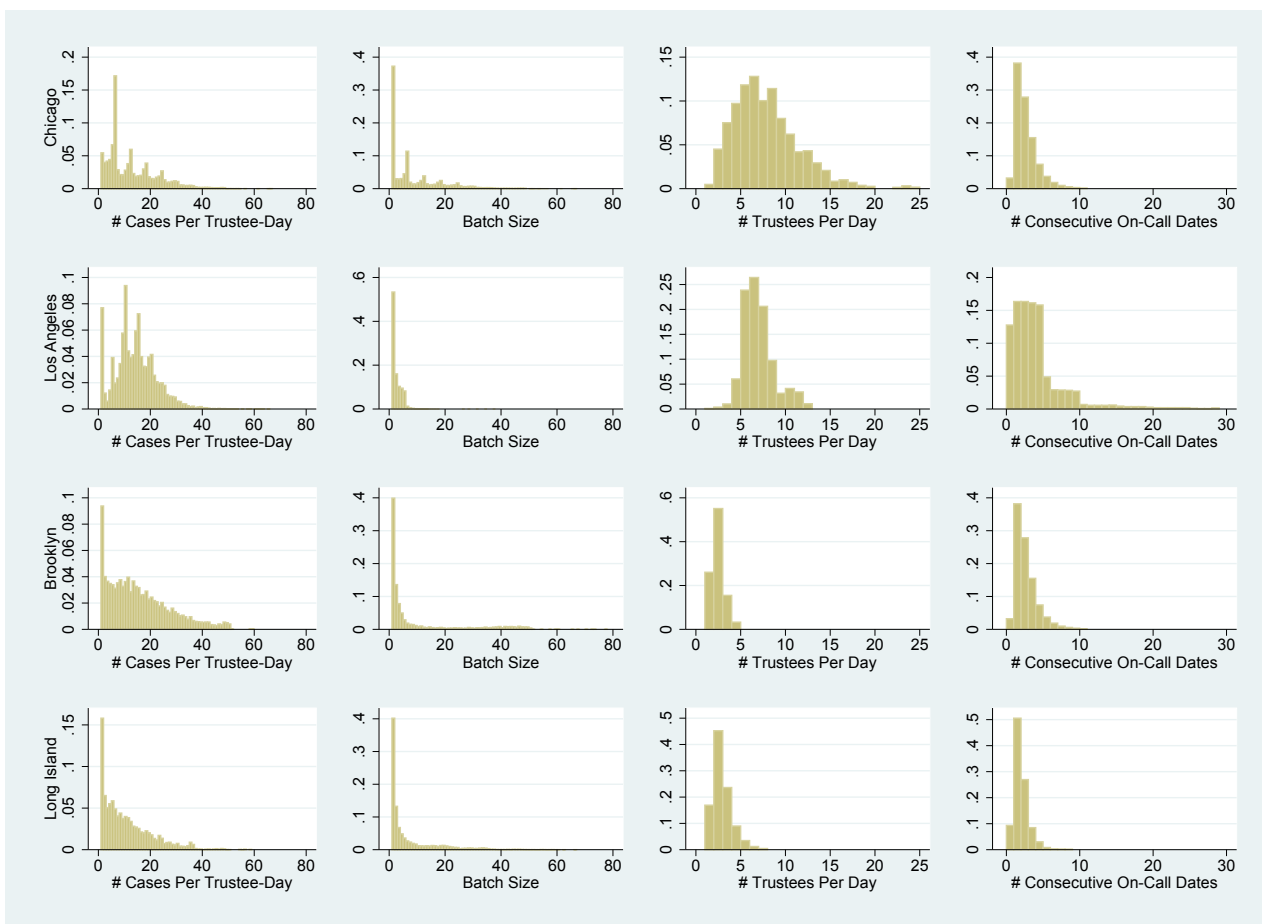
This Figure presents information on trustee actions, limited to pro se filers. The five bar charts show how frequently each trustee (a) holds a case open for more than 120 days; (b) adjourns a Section 341 meeting; (c) files at least 3 motions; (d) discovers assets to distribution to unsecured creditors; and (e) collects the debtor's tax refund to distribute to creditors, conditional on distributing to unsecured creditors.

FIGURE 7: Variation in Decision-Making Across Trustees, Full Sample



This Figure presents information on trustee actions, for all filers. The five bar charts show how frequently each trustee (a) holds a case open for more than 120 days; (b) adjourns a Section 341 meeting; (c) files at least 3 motions; (d) discovers assets to distribution to unsecured creditors; and (e) collects the debtor's tax refund to distribute to creditors, conditional on distributing to unsecured creditors.

FIGURE 8: Assignment Protocols, by Court



This Figure compares the trustee assignment protocols of four bankruptcy courts in Chicago, Los Angeles, Brooklyn, and Long Island. The first column of histograms shows the distribution of the number of cases per trustee per day. The second column shows the distribution of the number of cases in each assignment batch. The third column shows the distribution of the number of trustees on call during each day. The final column shows the distribution of the number of consecutive days that a trustee is on call.

Table 1: Summary Statistics

This Table contains summary statistics for our linked database of Chapter 7 cases filed or closed between October 2007 and October 2017 in Chicago. The identity of the trustee and law firm is from the "Cases Report" on the NDIL's ECF site. Debtor characteristics were scraped from filings on the NDIL's ECF site. The final assets distributed to creditors were scraped from filings on the NDIL's ECF site; a distribution case is one in which assets were distributed to creditors; a tax refund case is a distribution case in which one of the distribution assets is the debtor's tax refund.

(a) All Cases		
	Mean	Standard Deviation
Court Characteristics		
Cases per trustee per year	702.63	230.56
Trustees per year	38.38	2.33
Cases per law firm per year	25.49	139.70
Number of law firms per year	795.87	117.84
Debtor Characteristics		
Annual income	32,131.29	35,415.45
Assets	98,322.57	195,508.80
African American (%)	26.72	
Hispanic (%)	16.40	
Not African American, not Hispanic (%)	56.87	
Case Characteristics		
Pro se (%)	9.00	
Represented by top attorney (%)	13.50	
Represented by top 3 attorneys (%)	29.46	
Distribution case (%)	1.00	
Tax refund case (among distribution cases) (%)	32.0	
Number of Observations	181,709	

(b) Pro Se Cases		
	Mean	Standard Deviation
Debtor Characteristics		
Annual income	21,572.25	91,689.51
Assets	22,432.34	97,316.59
African American (%)	55.56	
Hispanic (%)	5.33	
Not African American, not Hispanic (%)	39.11	
Case Characteristics		
Distribution case (%)	0.74	
Tax refund case (among distribution cases) (%)	24.46	
Number of Observations	18,824	

Table 2: Evidence of Random Assignment Across Judges and Trustees

This table shows the results from a series of *F*-tests on differences in case characteristics across judges (or trustees). The null hypothesis for each test is that there is no difference among judges (or trustees). To obtain the *F*-statistics, we regress case characteristics, such as whether the consumer is African American, on a set of judge (or trustee) fixed effects plus month-year dummies. Each row of the table reports output from a separate regression, predicting the case characteristic in the first column. Columns 1 and 2 report *F*-tests (and significance levels) for the joint significance of the judge (or trustee) fixed effects. Columns 3 and 4 report the same *F*-tests but excluding debtors represented by an attorney.

	Full Sample		Pro Se Cases	
	(1)	(2)	(3)	(4)
	Judges	Trustees	Judges	Trustees
African American Debtor				
F stat	.49	1.95	.74	.66
p-value	.91	0	.70	.92
Hispanic Debtor				
F stat	.92	1.01	1.49	1
p-value	.52	.46	.13	.47
Filing Fee Paid in Full				
F stat	1.12	6.61	.49	1.12
p-value	.34	0	.91	.3
Original Asset Case				
F stat	1.27	2.18	1.37	.49
p-value	.24	0	.18	.99
Original Estimated Liabilities >50k				
F stat	1.39	1.73	1.04	1
p-value	.17	.01	.41	.46
Original Estimated Asset >50k				
F stat	1.32	3.86	.82	1.38
p-value	.2	0	.62	.08
Original Estimated ≥ 50 Creditors				
F stat	.91	1.59	.9	1.56
p-value	.53	.02	.54	.03

Table 3: Differences Across Trustees, Favored and Disfavored by Most Aggressive Attorneys

This Table compares intermediate and final case outcomes for NDIL cases filed by groups of attorneys that are assigned to certain trustees as compared to other trustees. To determine the group of disfavored trustees in the left column, we begin by determining the three attorneys with the highest mean absolute value of D_{aj} . D_{aj} is the difference between the number of matches between attorney a and trustee j and the expected number of matches, defined as the product of attorney a 's number of cases and trustee j 's number of cases divided by the total number of cases. We then select the three trustees with the smallest shares of these attorney's cases. Each Panel compares case outcomes with these trustees to other trustees for a subset of cases. Panel A subsets on cases by the three attorneys with high mean absolute value of D_{aj} . Panel B contains all attorneys but those included in Panel A. Panel C subsets on *pro se* cases.

	Disfavored Trustees	Other Trustees	P-Value
% Pro Se	13.84	9.71	0.00
Panel A: Attorneys with High Mean Absolute D Score			
% Adjourned	8.52	4.66	0.00
% Open More Than 4 Months	24.85	20.35	0.00
% with >2 Trustee Motions	7.10	2.74	0.00
% with Distribution	2.95	0.41	0.00
% Distribution Cases with Tax Refunds	26.47	5.83	0.00
% Debtors Who Subsequently File Again	8.08	7.24	0.28
Panel B: Other Attorneys			
% Adjourned	11.88	6.54	0.00
% Open More Than 4 Months	31.23	24.72	0.00
% with >2 Trustee Motions	11.77	4.59	0.00
% with Distribution	5.10	0.78	0.00
% Distribution Cases with Tax Refunds	20.98	8.72	0.00
% Debtors Who Subsequently File Again	9.32	9.17	0.56
Panel C: Pro Se			
% Adjourned	20.77	18.53	0.01
% Open More Than 3 Months	32.55	32.44	0.92
% with >2 Trustee Motions	15.19	11.94	0.00
% with Distribution	3.24	0.46	0.00
% Distribution Cases with Tax Refunds	44.16	2.74	0.00
% Debtors Who Subsequently File Again	33.81	34.41	0.57

Table 4: Random Assignment Protocols, by Court

This Table compares the assignment protocols across four courts. The Los Angeles column shows the average number of cases per trustee per day, the average batch size, the average number of trustees on call per day, and the average length of each on-call period. Each of the other three columns (Chicago, Brooklyn, and Long Island) not only shows the averages, but also gives a p-value from a two-sample t-test that compares the court with Los Angeles.

	Los Angeles	Chicago	Brooklyn	Long Island
# Cases per Trustee-Day	13.88	12.47 (0.00)	15.37 (0.00)	10.55 (0.00)
Batch Size (cases)	2.14	8.48 (0.00)	9.58 (0.00)	6.70 (0.00)
# Trustees Per Day	6.49	7.37 (0.00)	1.96 (0.00)	2.43 (0.00)
# Consecutive On-Call Dates	3.71	2.17 (0.00)	1.61 (0.00)	1.52 (0.00)

A APPENDIX

Table A.1: Test of Random Assignment - Abrams, Bertrand, and Mullainathan (2013) Method

This Table shows the p-values from a test of random assignment, following Abrams, Bertrand, and Mullainathan (2013). Each p-value is obtained in four steps: First, we remove trustees with fewer than 2,000 cases and de-mean case characteristics within the week of filing. Second, we rank trustees by their average value across all cases in the de-meaned case characteristic to be tested (i.e. log assets, log real property, etc.), and calculate the difference between the first and the thirty-first trustees for column 1, the third and the twenty-ninth trustees for column 2, etc. as well as the standard deviation across trustees. Third, we randomly re-assign trustees to cases within the week (so that each trustee keeps the same number of cases), re-rank the trustees, and re-calculate the differences. We run these simulations 300 times. Finally, we compute the empirical p-value of the actual spread as compared to the simulated spreads. The bottom panel is limited to *pro se* cases.

(a) All Cases						
	31 to 1	29 to 3	27 to 5	25 to 7	21 to 11	sd
Log Assets	0.00	0.00	0.00	0.00	0.00	0.00
Log Real Property	0.01	0.00	0.00	0.00	0.04	0.00
Log Personal Property	0.00	0.00	0.00	0.00	0.00	0.00
Log Liabilities	0.00	0.00	0.00	0.00	0.00	0.00
Log Secured Debt	0.01	0.00	0.00	0.00	0.07	0.00
Log Unsecured Priority	0.02	0.23	0.47	0.40	0.31	0.06
Log Unsecured Nonpriority	0.01	0.00	0.00	0.00	0.00	0.00
Log Dischargeable	0.00	0.00	0.00	0.00	0.00	0.00
Log Nondischargeable	0.06	0.14	0.02	0.21	0.37	0.05
Log Total Debt	0.00	0.00	0.00	0.00	0.00	0.00
Log Current Income	0.07	0.04	0.02	0.02	0.04	0.01
Log Average Income	0.00	0.00	0.00	0.00	0.16	0.00
Log Average Expenses	0.00	0.00	0.00	0.00	0.00	0.00

(b) Pro Se Cases						
	31 to 1	29 to 3	27 to 5	25 to 7	21 to 11	sd
Log Assets	0.58	0.16	0.19	0.33	0.46	0.34
Log Real Property	0.19	0.18	0.16	0.05	0.53	0.18
Log Personal Property	0.56	0.14	0.30	0.78	0.70	0.43
Log Liabilities	0.95	0.48	0.09	0.17	0.04	0.55
Log Secured Debt	0.38	0.10	0.36	0.43	0.55	0.31
Log Unsecured Priority	0.54	0.13	0.29	0.16	0.62	0.32
Log Unsecured Nonpriority	0.64	0.10	0.07	0.12	0.03	0.18
Log Dischargeable	0.80	0.47	0.43	0.15	0.11	0.55
Log Nondischargeable	0.68	0.35	0.14	0.25	0.70	0.47
Log Total Debt	0.71	0.40	0.26	0.05	0.12	0.38
Log Current Income	0.41	0.23	0.38	0.41	0.24	0.25
Log Average Income	0.10	0.08	0.05	0.10	0.90	0.08
Log Average Expenses	0.22	0.34	0.17	0.67	0.28	0.23