

Institutional Design of Misconduct Redress: Evidence from the Financial Ombudsman Service*

William Matcham[†]

May 26, 2026

Abstract

I study the design of institutions that adjudicate disputes between consumers and firms. The model I construct includes selection into challenges, pre-trial settlement, adjudication errors, and procedural fees. The fee schedule in the model matches my empirical setting—the UK consumer finance market—where lenders, but not consumers, are charged a fee irrespective of the case outcome. The fee structure creates rent-seeking opportunities through settlement by low-merit challengers. In all versions of the model I analyze, low-merit challengers go to trial only in exceptional circumstances, because lender fees enable pre-trial settlement. Hence, trial uphold rates overstate population-level rates of misconduct. I estimate the model using data scraped from trial documents and show that replacing lender case fees with a consumer deposit refunded upon victory saves innocent lenders up to \$1.3 billion per year but increases ombudsman net operating costs by around \$730 million.

Keywords: Misconduct, Litigation, Signaling, Institutions, Financial Redress

JEL Classification: C78, D82, G28, G50, K41

*I thank Mark Schankerman, Gregory Veramendi, and Melanie Lührmann for helpful discussions and comments.

[†]Department of Economics, Royal Holloway University of London, and STICERD, London School of Economics and Political Science, Houghton Street, London, UK. Email: william.matcham@rhul.ac.uk

1 Introduction

Misconduct—market actions that contravene legal, regulatory, or ethical norms—is a pervasive feature of several markets central to the economy. A primary example, and the focus of this paper, is the financial services industry.¹ Egan et al. (2024) documents that in some regions of the United States, as many as one in six financial advisers has a history of serious misconduct. The scale of U.S. financial advice misconduct is such that annual settlement costs in court cases for financial advisers in the period 2005–2015 alone totaled almost half a billion dollars. And in the United Kingdom, around \$50 billion (2026 USD) was refunded to consumers as a result of insurance mis-selling (FCA, 2020).

While the frequency and scale of financial misconduct are well documented, understanding how to reduce it is less explored, both in academic work and in policy. Most efforts have focused on two of the G20 Principles of Consumer Finance Protection: *financial literacy* (Lusardi and Mitchell, 2007; 2008; Lührmann et al., 2015; 2018) and *regulation* (Agarwal et al., 2014; Nelson, 2025; Matcham, 2025). This paper considers a third Principle: efficient *ex post* recourse mechanisms to resolve disputes. While consumers always have the option of pursuing redress through the courts, the costs of litigation are often prohibitive, and when no viable out-of-court mechanism is offered, consumers generally abandon their efforts to obtain a satisfactory resolution (Rutledge, 2010). Hence, institutions handling disputes have been created in many developed economies to provide a more accessible intermediate option.² Their efficient design is a key question for consumer financial protection, and is the overarching topic I study in this paper.

My primary contribution is to develop an empirical framework for analyzing the institutional design of ombudsman-style redress schemes, and to use it to quantify a central tradeoff between affordability and efficiency that arises whenever such institutions are funded by the parties they adjudicate. I show how imbalanced fee structures can lead to inefficiency through excessive challenges to innocent lenders.³ An overarching message is the tradeoff between affordability and efficiency in redress schemes. This tradeoff arises from rent-seeking opportunities that redress schemes can create when defendants fund the institution’s operation.

¹For leading court cases in other key markets, see [United States v. Holmes \(2018\)](#) for healthcare, [United States v. Blue Bell Creameries L.P. \(2020\)](#) for food safety, [United States v. Google LLC and YouTube, LLC \(2019\)](#) for tech and privacy, and [United States v. Volkswagen AG \(2017\)](#) for automobiles.

²This point extends beyond consumer finance, for example, to state-specific *small claims courts*, [housing \(Housing Ombudsman Service, 2024\)](#) and *patent* litigation (The U.S. Patent Trial and Appeal Board, [Love et al. \(2019\)](#)).

³By “inefficiency” in this context, I mean the presence of socially costly challenges by challengers with an invalid case.

My laboratory is the Financial Ombudsman Service (FOS), the UK institution that handles financial disputes and determines redress in the case of misconduct. I build, analyze, and estimate a model of this institution, and use the model to evaluate alternative designs of this institution. The model considers a dynamic game of redress capturing (i) selection into challenges by consumers, (ii) pre-trial settlement, and (iii) procedural fees. A consumer (the challenger) has a potential case against a financial services provider (abbreviated to a “lender”), which can be a high-merit (guilty lender) case or a low-merit (innocent lender) case. The challenger can initiate a complaint at a small cost. If the dispute does not resolve before FOS in the settlement stage, entering a FOS trial triggers a fee on the financial services provider. The “court” (FOS) is imperfect, meaning that there is a probability that guilty lenders are acquitted (type 1 error), and that innocent lenders are upheld (type 2 error). If a lender is found guilty, they pay the challenger damages.

My framework is sufficiently general to appeal to any redress institution combining (i) asymmetric fee allocation between disputants, (ii) imperfect adjudication, and (iii) the option to settle in the shadow of an institutional decision. These features are shared, in varying combinations, by financial dispute resolution schemes in Australia and other European economies, arbitration in the US, and non-financial ombudsman schemes in housing, energy, and telecoms (see Section 2 for further details).

The model delivers two primary insights. First, under complete information (i.e., both parties observe the merit), there is settlement in every case. Since trials impose a fixed lender fee *plus* damages on the financial services provider, there exists a mutually agreeable transfer interval, so even an innocent lender is willing to settle to avoid paying the fee. This result aligns with concerns about rent-seeking settlements and shows how a non-refundable defendant fee can generate transfers in low-merit disputes ([Financial Ombudsman Service, 2025](#)). While these settings clarify the inefficiencies arising from lender fees that are independent of the trial outcome, they cannot generate any trials in equilibrium, which is the outcome in up to 15% of FOS cases. As a result, I move to incomplete information settings.

The second insight is that when the lender cannot observe the merit, the model generates empirically relevant equilibria with occasional trials. Focusing on the case of the lender making the settlement offer, the lender faces a tradeoff: offer a low settlement and settle only in low-merit cases (and face trials against high-merit cases) or offer a high settlement transfer and settle with everyone. The lender pools (and hence always settles) when the proportion of guilty lenders exceeds a threshold that decreases in lender fees and increases in damages. A robust implication across equilibrium variants I analyze (including (i) the challenger making the offer, with separating and pooling possibilities, and (ii) the challenger not observing the merit) is that low-merit

challengers may *pursue* disputes opportunistically, but low-merit disputes only reach adjudication in exceptionally high damage cases, and never reach adjudication in my baseline incomplete information model. Consequently, observed uphold rates among adjudicated FOS decisions can substantially overstate the underlying incidence of misconduct in the population of disputes.

These theoretical results speak directly to current policy reforms. In 2025, FOS introduced fee changes that included fees for consumer representatives and partial refunds to lenders when acquitted. I show that the reform package cannot fully screen out low-merit challengers in the model: charging at the FOS stage can reduce the flow of cases into FOS, but it does so mainly by weakening the challenger’s threat point, making pre-trial settlement more attractive for both sides, more so than deterring opportunistic initiation. Separately, I show that removing lender fees entirely can go a long way towards removing opportunistic behavior. Still, it is unclear how such changes would be financed, since a large part of FOS funding comes from lender case fees.

I consider an alternative mechanism—a challenger deposit refundable upon a guilty finding—that can implement efficient screening (deterring low-merit pursuit while maintaining incentives for high-merit challengers), but at the cost of violating a core principle of the FOS model: being free at the point of use for *consumers*. Nevertheless, the FOS must be funded somehow, and if it is free for consumers, the burden falls on firms. This point elucidates a central tradeoff between affordability and efficiency in redress design, and motivates a quantitative question: how large are the inefficiencies generated by prioritizing free access? The empirical part of the paper estimates the structural model to answer this question.

For my empirical analysis, I collect multiple data sources. First, I scrape the near-universe of publicly available FOS decision documents (approximately 360,000). The dataset covers around 10% of total FOS cases. From the PDF decision letters, I extract case information (firm identity, product, outcome, etc.), ombudsman identity and gender, challenger gender, the overall sentiment of the decision letter, and (where feasible) the damages awarded in upheld cases. Descriptive analysis of the scraped data reveals an uphold rate near 33%, damages having a long right tail (mean of \$2,959; median \$340), and heterogeneity in uphold rates, mean damages, and sentiment across products and over time. I complement the scraped data with aggregate moments from the Financial Conduct Authority and Financial Ombudsman Service that allow me to estimate the model and conduct counterfactual analysis of alternative fee structures.

I estimate the prevalence of FOS errors, the time cost of consumer pursuit, and the underlying proportion of high-merit cases using simulated method of moments, matching on settlement rates, average settlement redress amounts, and rates at which consumers with issues pursue lenders. The model estimates imply a low level of misconduct and that trials are effective at acquitting

innocent lenders but relatively ineffective at convicting guilty lenders. The current system leads to approximately \$1.3bn per year in settlement payments by innocent lenders, and cases create a deficit (costs less case fees) of \$144m for the FOS.

I use the estimated model to conduct counterfactual analysis of alternative fee structures and institutional design. Removing FOS errors and eliminating lender case fees would eliminate innocent challenges but would increase FOS costs by \$227m, eliminate \$248m of revenues, and would need to be funded externally. A challenger deposit scheme would also save innocent lenders the entire \$1.3bn of payments per year. But it would increase FOS net operating costs by around \$730m per year due to the need to adjudicate more cases. This highlights the tradeoff between affordability and efficiency: while a challenger deposit scheme can eliminate opportunistic challenges, it would also increase FOS costs, and may deter meritorious claims. Funding the FOS by replacing case fees with a higher annual levy would improve efficiency and benefit lenders that do not engage in misconduct, at the expense of those who do.

Finally, it is worth noting that while the model parameters are calibrated to the UK FOS, the insights of the model and the underlying economic logic behind the reforms I study are *not* specific to the context, and are relevant for any dispute resolution mechanism for which the same economic logic applies. However, a natural caveat in interpreting my counterfactual analyses is that the underlying rate of misconduct is held fixed: lenders do not adjust their conduct in response to changes to fee structures or the quality of the FOS. Endogenizing misconduct in a more integrated models is a natural extension of the paper that I discuss in more detail in the conclusion.

Related literature

First, I contribute to the literature on models of litigation, settlement, and bargaining. Classic models emphasize how private information and litigation costs generate endogenous settlement and selection into trial, implying that observed trial outcomes need not be representative of the underlying population of disputes (Priest and Klein, 1984; Lee and Klerman, 2016). A large theoretical literature develops these ideas in more structured bargaining environments. Early contributions model litigation as a signaling game in which the act of going to trial conveys information about the plaintiff's type or the strength of the case (Salant, 1984; Bebchuk, 1984; Banks and Sobel, 1987; Sobel, 1989). A complementary approach studies settlement as a dynamic bargaining game in which parties make offers to avoid costly adjudication, with trial occurring when bargaining fails (Spier, 1992a;b; 1994; 1997; Fenn and Rickman, 1999). The model in this paper builds most directly on this latter tradition. My contribution relative to this literature is to develop a settlement-bargaining model tailored to the institutional features of ombudsman-style redress, and to estimate it for counterfactual analysis of institutional design.

Second, the paper relates to a rapidly growing literature documenting the scale and consequences of misconduct in financial markets. A sequence of papers ([Egan et al., 2019; 2022; 2024](#)) characterizes misconduct across financial services, emphasizing that misconduct can be pervasive, persistent, and unevenly distributed across firms and intermediaries. Recent work similarly highlights the economic relevance of misconduct and the importance of regulation over more expensive litigation in the courts ([Eliason et al., 2025; Annan, 2025](#)). The scale of misconduct documented in these papers motivates the need to understand how institutions should be designed to deliver redress and shape deterrence. One notable paper towards this endeavor is [Egan et al. \(2025\)](#), which estimates a model of financial arbitration in the U.S. They study the impact of the arbitrator selection mechanism using data on arbitration in the securities industry, and find that the effectiveness of the selection mechanism depends on the extent to which consumers are informed. Despite this important contribution, much of the misconduct literature still focuses on measurement, incidence, and consumer harm. Instead, I focus on institutional design: how the structure of redress—in particular, who bears the costs of initiating adjudication and how adjudicatory error affects incentives—changes equilibrium settlement behavior, adjudication rates, and ultimately the efficiency of the redress system.

There is recent work on the empirical analysis of screening institutions, particularly agencies that incentivize innovation. [Schankerman and Schuett \(2022\)](#) and [Freilich et al. \(2024\)](#) study the design of the U.S. patent system, including application, examination, and post-grant litigation, and [Matcham and Schankerman \(2025\)](#) studies the specifics of the design of the U.S. patent examination process. The present paper is similar in that I study the design of a public agency responsible for screening, but I focus on screening financial misconduct cases rather than patent applications.

Finally, there is a small literature on the effectiveness of financial complaint reporting schemes. Two recent papers study the effects of the Consumer Financial Protection Bureau (CFPB) complaints reporting system in the United States ([Dou et al., 2024; Vaishnav et al., 2024](#)). These papers emphasize that complaint reporting can discipline firms and generate information for regulators and consumers even without a centralized adjudicator. The FOS differs from the CFPB complaints process in that it provides a structured investigation and can issue determinations that bind firms. This institutional distinction creates an additional strategic margin: the threat of referral induces settlement bargaining in the shadow of adjudication.

2 Context

Before the model, I provide an overview of the key features of the UK Financial Ombudsman Service (FOS) and the recent changes to its fee structure. A reader familiar with the context could skip this section and move straight to the model. The FOS is an organization that resolves disputes between consumers and financial service providers in the UK. Established in 2001, it is an independent public body intending to provide an impartial alternative to the courts for consumers (and eligible small businesses) who have disputes with regulated financial services firms. In contrast to court proceedings, the process is designed to be less involved and faster but more accessible: consumers do not need legal representation, and there is no filing fee for complainants.

Although the FOS is institutionally independent in its casework, it sits within the UK’s broader regulatory architecture. The financial regulator (the Financial Conduct Authority, FCA) plays a central role in governance and funding; for example, the FCA sets (or approves) the overall budget and appoints the FOS board. While the FCA cannot directly intervene in individual decisions, the two bodies share statutory duties to coordinate on matters with “wider implications” for the financial system and consumer protection. This separation of adjudication from regulatory enforcement is a defining feature of the UK ombudsman model: the FOS resolves individual disputes, while the regulator oversees market-wide conduct. The FCA is actively seeking to modernize the FOS ([FCA and FOS, 2025](#)), but the academic literature offers limited guidance on how to do so. The present paper offers such guidance through an empirical model of the Ombudsman.

The timeline of the complaints process is as follows. A consumer (or their representative) must first complain directly to the business. The business then has eight weeks to investigate and provide a final response. If the consumer is dissatisfied with the firm’s response—or if the firm does not respond within the eight-week window—the consumer may escalate the dispute to the FOS. There is a time limit for doing so: after receiving the firm’s final response, the consumer typically has six months to refer the complaint. Once referred, the FOS first checks whether the complaint falls within its jurisdiction. If so, the FOS gathers evidence from both sides, investigates the dispute, and bills the lender for the case (currently approximately \$880 to be paid to the FOS at the end of the case, *irrespective* of outcome). The FOS then issues a provisional, and then (if needed) a final decision. If the consumer accepts their decision, it becomes binding on the firm. If the consumer rejects the FOS final decision, they retain the option to pursue the

matter through the courts.⁴

The FOS can require firms to provide redress intended to put the consumer back into the position they would have been in absent misconduct, and, in some cases, to provide compensation for distress and inconvenience caused. The maximum award has increased over time and is currently about \$600,000 for non-pension cases.

Other economies have similar ombudsman processes. Several jurisdictions use specialized alternative dispute resolution systems for financial complaints, including Australia (AFCA) and Ireland (FSPO). Some systems deliver legally binding decisions, while others emphasize mediation or informal resolution. The United States has no direct equivalent to the UK FOS. Instead, the closest analogue is an administrative complaints process through the Consumer Financial Protection Bureau (CFPB), which collects complaints and facilitates responses but does not provide the same adjudicative mechanism or binding determinations. Similarly, the Financial Consumer Agency of Canada is also not an ombudsman but a complaint handler. In these cases, if the lender rejects a complaint, the consumer has limited options but to take them to court or to arbitration.⁵

Funding and Fees A key institutional feature—and the one most important in this paper—is how the FOS is funded. The system is free at the point of use for consumers. It is financed by regulated firms through two channels: (i) a compulsory levy on regulated firms, and (ii) per-case charges when complaints are referred to the FOS and accepted into its process. Hence, disputes impose costs on defendants, even before any decision on the merits.

Lender case fees have increased over time roughly in line with inflation, from around \$500 in 2003 to \$1050 in 2021. An important change to the fee structure was made in April 2025. First, the lender fee was reduced to \$880 and further reduced by around \$250 if the case against them failed. Second, in response to vast increases in cases brought forward by professional representatives (PR)—typically claims management companies, abbreviated to CMC—any PR must pay a fee of \$350 for each subsequent case after its tenth each year. However, if the complaint is successful, the PR is refunded approximately \$250 in credit. However, consumers continued to face no ombudsman fees for bringing forward a case themselves. In Section 4, I examine the effect of this

⁴This asymmetry is important for the model that follows. The consumer’s post-FOS outside option (litigation) is incorporated implicitly through the trial payoff, and is empirically small given the costs of UK civil litigation relative to typical FOS damages. The lender, by contrast, has no merits-based right of appeal. The only option is a costly judicial review, available only on narrow procedural grounds and rarely successful.

⁵These agencies do monitor the number of complaints for specific matters to find widespread issues, but idiosyncratic issues are less likely to be addressed.

fee change within the model I construct.

3 A Model of Financial Redress

3.1 Setup

I model the financial redress process as a dynamic game between a potential ChALLENGER (consumer or a PR), who is the appellant and denoted by C , and a financial services provider, who is the defendant and denoted by L (e.g., a Lender).⁶

The challenger is endowed with a potential redress case of merit $\theta \in \{G, I\}$ (either the lender is Guilty or Innocent) chosen by Nature, with the proportion of guilty cases in the population being equal to $\pi = \mathbb{P}(\theta = G) \in (0, 1)$. When $\theta = G$, the case is said to be of high-merit; $\theta = I$ corresponds to low-merit cases. The timing, as in Figure 1, is as follows:

Stage 1. Case Initiation: The challenger's first choice is whether to present the case to the lender. Not presenting the case yields a payoff normalized to zero and ends the game. Otherwise, they prepare the case at the administrative cost of $K > 0$, and pursue the lender.

Stage 2. Pre-FOS Settlement: After presenting the case to the lender, but before going to the FOS, there is a potential for pre-trial settlement, in which the lender pays the challenger the sum of T to drop the case. At this point, if pre-trial settlement is agreed, the transfer is made, and the game ends. Otherwise, the case goes to a FOS trial.


Stage 3. FOS Trial: The FOS bills the lender for f_L to be paid at the conclusion of the FOS case, irrespective of the outcome and the value of θ . In the FOS trial, the ombudsman comes to a decision about the merit of the case. If the FOS finds the lender guilty, the lender must pay the challenger damages of $B > 0$.⁷ To allow for the possibility of type 1 and 2 errors in FOS decisions, the probability that the FOS finds type θ to be guilty is $e_\theta \in [0, 1]$.⁸ The FOS are more likely to uphold a guilty lender than an innocent one, i.e., $e_I < e_G$. A perfect FOS would have

⁶I do not distinguish in the basic model between a when a consumer makes the challenge and a PR represents them. The choice to use a PR can be thought of as an increase in K and e_θ , explained later.

⁷Since the consumer can escalate the case to the court if they are not happy with the decision, B can be interpreted as the expected winnings inclusive of the expected returns from escalating to a court case. As I noted in Section 2, the lender essentially cannot challenge the decision of the FOS in court unless they go for a judicial review, which is a costly and rarely successful option.

⁸As referenced in footnote 6, while the model does not endogenize the use or role of representative, the motive for using one is that it would increase the chance of winning the case, but cost the consumer more, so the use of a PR can be thought of as a higher K and higher e_θ .

FIGURE 1. Model Timing



$e_I = 0$ and $e_G = 1$.

The parameters π, e_G, e_I, K, f_L , and B are common knowledge.⁹ In the main text, I distinguish the cases of θ being common knowledge (Section 3.2) and private information of the challenger (Section 3.3). Finally, while the parameters are labeled with FOS-specific names (lender, FOS, etc.), the structure maps to a broader class of institutions involving a defendant and challenger, an adjudicator with errors, and potentially asymmetric fee structures.

3.2 Complete Information Benchmark

To start, suppose that the realization of θ is known to challenger and lender. In this case, the subgame perfect equilibrium is obtained via backward induction. The final decision is whether to settle. For each θ , at the point of pre-trial settlement, if settlement is not agreed, the lender stands to lose $f_L + e_\theta B$ and the challenger stands to gain $e_\theta B$. Hence, both parties would accept a transfer $T \in [e_\theta B, f_L + e_\theta B]$ from lender to challenger. This region covers Nash Bargaining solutions across the full spectrum of bargaining weights, and the extreme cases of the lender making a take it or leave it (TIOLI) offer (in which they would choose the lower boundary $T_L = e_\theta B$) and the challenger making a TIOLI offer (in which they would choose the upper boundary, $T_C = f_L + e_\theta B$).

Going back one step, the challenger pursues the lender if $T > K$, which is ensured if $e_I B > K$. The administrative costs of complaining, K , are small in practice, consistent with the large volume of complaints for non-trivial B values (see Section 5).

The equilibrium path can involve any T in the region above, and has the challenger bringing forward a case if $T > K$. The main observation is that the lender always settles on the equilibrium path, even when they are innocent. Because the lender must pay f_L irrespective of the case's merit, *even if* $e_I = 0$, they are still better off settling when they are innocent than going to trial and winning.

⁹Alternatively, though, one can think of B as not known until final outcome, independent of all other model variables, and that B in the model description represents $\mathbb{E}(B)$.

In the complete information case, the policy changes to remedy the equilibrium are straightforward. If $e_I = 0$ (i.e. the FOS never finds innocent lenders guilty), then the simple change to a system in which f_L is only charged if the lender is guilty suffices. In that case, provided $e_G B > K > 0$, challengers never present low-merit cases to lenders and always present high-merit cases to lenders, which are subsequently settled. The threat of FOS trials is enough to make FOS trials unnecessary on the equilibrium path.

When $e_I > 0$, it is not enough to remove f_L for innocent lenders alone, as challengers with low-merit cases but high enough damages can use their chance of getting a lucky decision in a trial to force a settlement. Hence, a viable policy effort is to reduce e_I . Further, if we supplement the above remedy with a change requiring challengers to pay a sufficiently large f_C when the lender is found to be innocent, then we obtain the same outcome as in the case of $e_I = 0$. Though it is understood that the FOS does not want to impose fees on consumers, the simple model shows that it is hard to stop rent-seeking challenges by low-merit consumers unless there is some fee disciplining their behavior.

3.3 Incomplete Information

The previous benchmark formalizes the argument made by opponents of the FOS fee structure. In the benchmark, all low-merit challengers file cases and achieve a pre-trial settlement, despite it being common knowledge that the lender is innocent.

The disadvantage of the simple model is that guilty lenders also never reach trial, so there are no trials in equilibrium. This feature makes the complete information model unsuitable for empirical implementation, since the data indicate that up to 15% of cases go to trial in some years. To generate FOS trials in equilibrium, I move to models in which the merit of the case is not known to one of the two agents.

In the main text, I cover the case in which the challenger but not the lender knows the merit. I discuss why I focus on this particular case in Section 3.5, and cover other cases in Appendix B. Once I include incomplete information, I have to take a stance on which of the challenger and lender makes a settlement offer. When θ is not known to the lender, it is not possible to analyze general models of pre-trial settlement in which the value of the transfer is determined by the model. Indeed, the cases modeled in the literature either require that the lender or the challenger make a settlement offer to the other party. I focus in the main text on the equilibrium with the lender proposing the settlement offer, because this is (i) most realistic (what happens in practice is that the lender proposes an amount to the challenger after a complaint), (ii) simpler to analyze, and (iii) has an intuitive equilibrium generating trials with no need for equilibrium

refinements. In Appendix B, I analyze pooling and separating equilibria in which the challenger makes the settlement offer. When the challenger makes a settlement offer, the model becomes a signaling game in the spirit of the [Spence \(1973\)](#) model.¹⁰

3.4 Lender Ignorance and Lender Settlement Offer

In this case, the lender chooses a settlement payment T to minimize its expected cost, taking into account how different offers affect which types of challengers accept settlement versus proceed to trial. In this case, the equilibrium is in similar spirit to that of a Spier equilibrium (e.g., [Spier, 1997](#)).

3.4.1 Small Cost of Pursuit

I start with the case of $K < e_I B$, so that both types pursue the lender in the first step, for all T . Hence, the lender cannot infer on type from the filing decision. This case provides the main intuition for this model setup. I move to the more involved case of $K > e_I B$ in Section 3.4.2.

In this environment, one candidate is to offer $T = e_I B$. With this offer, low-merit challengers (innocent cases) are exactly indifferent between accepting the settlement and proceeding to trial, while high-merit challengers prefer to proceed to trial. Hence, there is settlement only with low-types, and the lender's expected cost under this offer is therefore

$$EC_I := (1 - \pi) e_I B + \pi (f_L + e_G B) = \underbrace{(1 - \pi) e_I B + \pi e_G B}_{E(\text{Damage})} + \pi f_L,$$

which reflects settlement payments to innocent cases plus the trial cost borne when the lender is guilty.¹¹

Alternatively, the lender can offer $T = e_G B$, which is sufficiently generous that both types of challengers accept settlement. In that case, no dispute proceeds to trial and the lender's cost is $EC_G = e_G B$.¹² Comparing EC_I and EC_G , the lender prefers to settle with everyone (choose $T_G = e_G B$) whenever

$$e_G B < (1 - \pi) e_I B + \pi (f_L + e_G B).$$

¹⁰[Banks and Sobel \(1987\)](#) contains an example of a signaling game like a court case, further analyzed in [Salant \(1984\)](#) and [Sobel \(1989\)](#).

¹¹The lender will never offer $T < e_I B$ as both type of challengers will reject that offer, so the expected cost will be $E(\text{Damage}) + f_L > EC_I$.

¹²The lender will never offer $T \in (e_I B, e_G B)$ as that will cost more than $e_I B$ but not induce a different type of challenger than $e_I B$. Similarly, the lender will never offer $T > e_G B$ as that will still be accepted by both type of challengers but cost more than $e_G B$.

Rearranging yields a simple threshold condition:

$$T = e_G B \iff \pi > \pi^* \equiv \frac{\Delta_e B}{f_L + \Delta_e B}, \quad (1)$$

where $\Delta_e := e_G - e_I$.¹³ Intuitively, if the lender is sufficiently likely to be guilty (high π), it becomes optimal to avoid the expected cost of trial—which includes both the adjudication fee f_L and the expected liability—by making a settlement offer generous enough to eliminate adjudication entirely. In what follows, I call the cases $\pi > \pi^*$ Case 1, and $\pi \leq \pi^*$ Case 2.

A key implication of this structure is that *low-merit challengers never reach trial*. Regardless of π , the lender can always make an offer that is accepted by innocent types at weakly lower cost than allowing those cases to proceed to adjudication. Thus, any disputes that proceed to a final decision are disproportionately those in which the lender is guilty (type G).

3.4.2 Intermediate and Large Cost of Pursuit

Now, I move to the case of $K \in [e_I B, e_G B]$. Note that $K > e_G B$ is not interesting as neither types pursue the lender in equilibrium in that case. The region $[e_I B, e_G B]$ is interesting because low-merit challengers (type I) are no longer guaranteed to pursue the lender: pursuing is potentially unprofitable unless settlement terms are favorable. As a result, the lender's settlement offer can affect not only whether cases settle or proceed to trial, but also whether low-merit challengers enter the dispute at all.

Case 1: $\pi > \pi^*$. When the lender believes it is sufficiently likely to be guilty, it is optimal to avoid trial risk and the case fee by settling with everyone. In this region, the equilibrium coincides with the case $K < e_I B$ analyzed in the main text: the lender sets $T = e_G B$, and both types accept.

Case 2: $\pi \leq \pi^*$. When π is below the cutoff, the lender would prefer to avoid paying $e_G B$ to all challengers. However, the option $T = e_I B$ is no longer optimal in pure strategies in the region $K \in [e_I B, e_G B]$.

To see why, note that if the lender offered $T = e_I B$, then only low-merit challengers would be exactly indifferent between settling and going to trial *conditional on having pursued*, but because $K \geq e_I B$, low-merit challengers may prefer not to pursue in the first place. In the extreme, if low-merit challengers never pursue, the lender's offer $T = e_I B$ would only ever be faced by type G challengers, who would reject and proceed to trial, making this offer suboptimal for the lender.

¹³If $\pi = \pi^*$, the lender is indifferent, and any $T \in [e_I B, e_G B]$ is optimal.

Instead, the equilibrium in this region takes a mixed-strategy form in which both the lender and low-merit challengers randomize. In the mixed-strategy equilibrium, the high-merit type $\theta = G$ always pursues the lender and the low-merit type $\theta = I$ pursues with probability $\lambda \in [0, 1]$. The lender mixes between two settlement offers: $T = e_I B$ with probability σ and $T = e_G B$ with probability $1 - \sigma$.

Acceptance behavior follows the standard cutoff logic: $\theta = G$ accepts if $T \geq e_G B$, while $\theta = I$ accepts if $T \geq e_I B$. Thus, $T = e_G B$ induces settlement with both types, while $T = e_I B$ induces settlement with I but leads G types to proceed to trial. The mixing probabilities (λ, σ) are pinned down by indifference conditions.

Determination of λ . The lender chooses a mixed strategy only if it is indifferent between its two candidate offers. If the lender offers $T = e_G B$, everyone settles, so its cost is $(1 - \pi)\lambda e_G B + \pi e_G B$. If the lender offers $T = e_I B$, then I types (when they pursue) settle and G types go to trial. Because I pursues with probability λ , the lender's expected cost from the low offer is

$$\underbrace{(1 - \pi)\lambda e_I B}_{\text{settle with pursuing } I} + \underbrace{\pi(f_L + e_G B)}_{\text{trial against } G}.$$

Indifference implies

$$(1 - \pi)\lambda e_G B + \pi e_G B = (1 - \pi)\lambda e_I B + \pi(f_L + e_G B),$$

which yields

$$\lambda = \frac{\pi}{1 - \pi} \frac{f_L}{\Delta_e B}.$$

The value of λ is in $[0, 1]$ if and only if $\pi f_L \leq (1 - \pi)\Delta_e B$, which is ensured by the definition of Case 2, which has $\pi < \pi^*$.

Determination of σ . Low-merit challengers mix only if they are indifferent between pursuing and not pursuing. If a low-merit challenger does not pursue, its payoff is 0. If it pursues, it pays cost K , and then receives a settlement of $e_I B$ if the lender makes the low offer (probability σ), and receives $e_G B$ if the lender makes the high offer (probability $1 - \sigma$). Thus, the expected payoff from pursuing is $\sigma e_I B + (1 - \sigma)e_G B - K$. Indifference implies $\sigma e_I B + (1 - \sigma)e_G B - K = 0$, which rearranges to

$$\sigma = \frac{e_G B - K}{\Delta_e B}.$$

Because $K \in [e_I B, e_G B]$, this probability lies in $[0, 1]$, ensuring the mixed strategy is well-defined.

This equilibrium generates trials with strictly positive probability. Trials occur precisely when the challenger is guilty ($\theta = G$ pursues surely) and the lender draws the low settlement offer

$T = e_I B$ (probability σ), which G rejects. Low-merit types enter only probabilistically (through λ), and when they do enter, they always settle.

The mixed-strategy structure therefore preserves the key qualitative prediction from the baseline case: low-merit cases are disproportionately resolved through settlement (or non-entry), while trials are concentrated among guilty cases.

3.5 Equilibrium Implications Under the Old Fee Structure

The analysis above provides my primary equilibrium characterization under the baseline fee regime. There are four alternative modeling environments: (i) both parties having incomplete information on θ (analyzed in Appendix Section B.1); (ii) lender incomplete information with challenger offers, which can generate separating or pooling equilibria (analyzed in Appendix Sections B.2.1 and B.2.2); (iii) challenger incomplete information with challenger offer (analyzed in Appendix Section B.3.1); and (iv) challenger incomplete information with lender offer (analyzed in Appendix Sections B.3.2 and B.3.3).

While exact conditions for the equilibrium and the equilibrium path itself differ across informational assumptions and choice of party making settlement offer, one result is common to every one of the models analyzed: under reasonable parameterizations and aside from cases with exceptional damages, the FOS is unlikely to judge on low-merit cases in equilibrium, in other words, only guilty lenders will reach a FOS trial. Hence, because innocent cases are disproportionately screened out through settlement, the uphold rate among adjudicated decisions overstates the underlying fraction of guilty cases in the population. In fact, when only guilty cases proceed to a trial, the uphold rate in observed decisions is exactly e_G .

There are two points worth noting regarding my choice of baseline in the imperfect information case (i.e., the lender not knowing θ and the lender making a settlement offer). First, on a practical level, each customer could know their personal circumstances better than a lender, when a lender is dealing with the individual circumstances of potentially millions of customers and investigating specific cases is costly. Alternatively, one can interpret the lender's ignorance as not knowing how the FOS will *judge* the case (mistakes aside), or not knowing the details of the relevant regulatory requirements, rather than not knowing what happened in each individual case. Second, and from a modeling perspective, a main aim of the following section will be to see whether policies can deter low-merit challengers from applying while maintaining high-merit challenges. When challengers do not know their type, it is not possible to think of a way to deter low-merit challengers, so such a task becomes futile in the case of ignorant challengers.

Before moving to the theoretical analysis of reforms, I note that the model I construct treats

each case separately. This choice is fine for idiosyncratic consumer issues, but makes the model less appropriate for mass redress events such as the recent Payment Protection Insurance and Motor Finance redress events in the UK.

4 Theoretical Analysis of Model Reforms

This section uses the equilibrium described in the incomplete information model of Section 3.4 to analyze the effects of reforms on equilibrium outcomes theoretically. I provide intuition for the results in the main text, with formal details in Appendix C.

4.1 Removing FOS Errors and Lender Case Fees

First, consider the impact of just removing FOS errors ($e_I = 0$, $e_G = 1$) but keeping lender case fees. Because $e_I = 0$, $K > e_I B$ in all cases, so the equilibrium is as in Section 3.4.2. Otherwise, there is no difference: low types will either always pursue (Case 1) or pursue with non-negative probability (Case 2). Because of lender case fees, making the FOS free of error is insufficient to remove all low-merit challenges.

Next, consider the impact of removing lender case fees but maintaining FOS errors. With $f_L = 0$, $\pi^* = 1$ so we are always in Case 2. When $K \in [e_I B, e_G B]$, we have $\lambda = 0$, so that low-merit challengers pursue with probability 0 in this case. But when $K < e_I B$, we still see low-merit challengers pursue and settle at the low offer (which is always made) in this case. Hence, removing lender case fees with an imperfect FOS is still insufficient to eliminate low-merit challengers, as FOS mistakes make low-merit challenges worthwhile on average provided the costs of pursuit are not too high.

It is only when FOS errors are removed *and* lender case fees removed that low-merit challengers are eliminated entirely. Because $e_I = 0$, we never have $K < e_I B$, and because $\pi^* = 1$, we are always in Case 2 with $\lambda = 0$. Lenders still randomize to maintain the equilibrium, and high-merit types always pursue. While this could be described as an “efficient” outcome, in practice, removing lender case fees would increase the FOS deficit, and would necessitate an increase in alternative sources of funding for the FOS. Further, it is not immediate how much investment would be required to remove all FOS errors, if this is possible at all.

4.2 April 2025 Fee Reform

Next, I analyze the impact of the April 2025 change to the fee structure in the context of the model. Motivated by large increases of the volume of complaints, the reform introduced three economically important changes to the pre-existing fee structure.

First, a new fee for challengers' representatives to refer a case to the FOS, denoted f_C , was introduced.¹⁴ Second, if the ombudsman upholds the complaint (i.e., the lender is found liable), the representative receives a partial refund of some fraction of f_C . Third, the reform also adjusts the fee borne by the lender, denoted f_L , by providing a partial refund when the complaint is not upheld. This reduces the expected cost to lenders from facing non-meritorious or unsuccessful complaints, and therefore weakens the settlement pressure generated by low-merit challenges.

The natural question is whether introducing a challenger-side fee f_C for each FOS trial can *screen out* challengers of low merit, in the sense of preventing weak disputes from being pursued to the FOS. As I show in Appendix Section C.1, a challenger fee reduces the volume of disputes that reach final adjudication, but does not fully screen out low-merit cases.

The intuition is as follows. In the baseline environment, escalation to the FOS imposes a fixed expected cost on the lender (through f_L and the risk of an adverse decision), while the challenger can escalate at little or no cost. This gives even low-merit challengers leverage: they can threaten escalation in order to obtain a settlement payment that compensates the lender for avoiding the case fee and adjudication risk.

Introducing f_C raises the challenger's cost of escalation, which weakens the threat point in bargaining. This *slackens* the settlement condition: for a given underlying merit, the set of settlements that both parties prefer to adjudication expands. As a result, fewer cases proceed to the FOS, but the reason is that both parties become more willing to settle *before* the FOS. The fee can make escalation privately unprofitable for the challenger, and thus reduce adjudication volume, but it will not, in general, implement an entirely clean separation of high- and low-merit disputes. Since my model cannot distinguish between consumer challenges with and without a representative, I do not analyze this reform in more detail in the empirical component of the paper.

4.3 Refundable Challenger Deposits

The model also clarifies what *could* screen out low-merit challengers. Consider an alternative policy in which the challenger must pay an upfront deposit to pursue a lender, and the deposit is refunded if the lender is found liable at adjudication. In the model, such a deposit creates a direct screening mechanism: low-merit challengers have a lower probability of refund and therefore face a high expected cost from pursuing the dispute, while high-merit challengers expect to recover

¹⁴Without professional representatives as an explicit entity in the model, I consider this as a fee for the combined challenger and professional representative unit

the deposit and therefore are less deterred. As I show in Appendix Section C.2, provided $K < B$, there is always a deposit fee for which low-merit disputes are not pursued, but high-merit ones are.

However, while a refundable deposit can improve efficiency by reducing opportunistic escalation, it violates a central design principle of the FOS—that the process is free at the point of use for consumers. An upfront deposit may deter not only low-merit cases but also meritorious complaints brought by liquidity-constrained consumers, and it may reduce access to justice even when claims are socially valuable. This highlights the core policy tradeoff emphasized throughout the paper: affordability versus efficiency. Screening instruments that eliminate low-merit challenges often require imposing meaningful expected costs on challengers, but doing so undermines the objective of providing free and accessible redress. This tradeoff motivates an empirical implementation of the theoretical model to quantify the costs on low-merit lenders of the zero-fee policy for challengers; the remainder of the paper is devoted to this task.

5 Empirical Analysis

5.1 Data Sources

I exploit two complementary sources of information on the UK consumer financial dispute resolution system. The first source consists of aggregate complaints, settlements and case-flow statistics published by the FCA and the FOS, respectively. The second source of data comes from the universe of publicly available FOS final decision letters, which I collect via web scraping and parse to construct a case-level dataset containing case outcomes and characteristics.

Consumer Complaints and FOS Cases Firms are legally obliged to report complaints to the FCA, and the FCA publishes aggregate statistics on the total number of complaints made to UK financial service providers every six months. In the first half of 2025, financial services firms received 1.85 million complaints. The FCA also publishes aggregate statistics on settlement amounts across these complaints. Separately, the FOS publishes the total number of cases it receives each quarter. In the first half of 2025, the FOS received approximately 156,000 new complaints. I extract moments relating to complaint probabilities, trial probabilities, and average settlements from these datasets to estimate the model. I describe the precise moments I use in Section 6.

Public FOS Final Decisions The core case-level dataset in the paper comes from publicly available FOS final decisions. Around 10% of FOS cases are made public on the FOS website.¹⁵ I scrape the universe of publicly available final decisions (approximately 360,000 cases). The first step extracts a small set of structured fields displayed on the FOS public decisions portal, including (i) the decision reference ID, (ii) the decision date, (iii) the identity of the defendant, (iv) the case outcome (e.g., upheld), and (v) the broad product category (banking, credit, insurance, investments, etc.). Figure A.1a provides an example of the information available on the public listing page.

Then, each public listing links to a decision letter. I download the decision documents and extract the text in order to construct additional case characteristics that are not available in the structured listing. Figure A.1b shows an example decision document. From the decision text, I parse the ombudsman’s name (and a gender proxy inferred from first names) and the honorifics of challenger(s) (e.g., “Mr”/“Mrs”) as a proxy for perceived challenger gender. I also use standard methods in text sentiment analysis (Rinker, 2021) and readability analysis (Flesch, 1948) to score the tone and complexity of the correspondence. Finally, in upheld cases, the document reports redress and compensation amounts, which I interpret as observed damages (the outcome variable B in the model). Figure A.1b illustrates the parts of the document from which this information is extracted.

5.1.1 Damages Measurement and Limitations

While decision documents allow the construction of a damages measure, two limitations should be noted. First, damages in public cases are reported only when a case is upheld (or upheld in part), because unsuccessful decisions typically do not disclose the monetary stakes of the dispute. Second, automated extraction is most reliable for cases in which monetary relief is stated as a lump-sum amount (e.g., “pay £100 in compensation”). Some decisions instead specify remedies as a percentage of an underlying amount or as non-monetary corrective actions, which are harder to convert into a single monetary measure. As a result, the extracted damages measure is best interpreted as a lower bound on total compensation when percentage-based awards are present.

To assess robustness to the latter issue, I take two steps. First, I attempt to extract damages

¹⁵When a case goes to FOS, if the assigned investigator can handle the case without issue, the case is private and FOS do not publish the outcome. If the assigned investigator is unsure, or either party disagrees with the outcome the case goes to an ombudsman and the decision is made public. It is worth noting that ombudsmen agree with the original investigator on the outcome and damages in approximately 90% of cases, suggesting that public cases are representative of the full population of cases considered by the FOS.

using text parsing techniques and also by prompting a large language model (GPT-5 API) to read the decision letter and estimate total monetary redress as a single sum. Through manual checks of over 100 documents, I find that monetary sums from the two methods are identical in most cases, and where they differ, neither approach dominates. For the analysis that follows, I use the values reported by the large language model, as it is less susceptible to large mistakes (estimates are similar from both approaches, however). Second, I conduct robustness checks restricting attention to products where monetary damages are most likely to be explicitly stated, and I separately exclude cases in which the decision text contains a % symbol (which is a strong indicator that redress is expressed proportionally rather than as a lump sum).

5.2 Summary Statistics and Descriptive Analysis

This section presents summary statistics and descriptive facts about FOS outcomes and awards. The aim is to give the reader a general overview of the variables I observe in the data, and to highlight the heterogeneity the model aims to capture.

5.2.1 Summary Statistics

Table 1 reports summary statistics for the main sample of final ombudsman decisions. The sample contains just under 351,000 decisions. The main outcome variables are an indicator for whether the complaint is upheld, which occurs in 33% of cases, and damages when the case is upheld.

The damages distribution is highly skewed: mean damages are around \$2,959 with a standard deviation of \$16,575, while the median upheld award is \$340. The gap between mean and median, and the long-tailed nature of award outcomes, is considered in more detail in Section 5.2.2.

The distribution of sentiment is well-approximated by a normal distribution centered at 0 with a standard deviation of 0.04, implying a neutral tone on average. However, as expected, sentiment towards the consumer is positive when the case is upheld and negative when the lender is acquitted. The Flesch variable is the [Flesch \(1948\)](#) readability measure, where a higher value indicates that the document is easier to read. This measure is well approximated by a normal distribution centered at 60 with a standard deviation of 6. The value of 60 sits right on the border of what many guides call “standard/plain English” (60–70) versus “fairly difficult” (50–60) to read.

The composition of cases varies across broad product categories. Insurance accounts for 46% of decisions, followed by Banking (37%), then Investment/Pension/Annuities (10%), and finally Consumer Credit (9%). These differences are relevant because the distribution of underlying consumer costs of misconduct is likely to be highly product-dependent, as I explore in detail in Section 5.2.3. Finally, I observe characteristics of challengers and ombudsmen. There is an

TABLE 1. Summary Statistics

Variable	Mean	S.D.
Upheld	0.33	0.47
Damages Upheld (\$)	2,959.67	16,575.16
Document Sentiment	0.00	0.04
Flesch Readability	59.54	6.38
Banking	0.37	0.48
Consumer Credit	0.09	0.29
Insurance	0.46	0.50
Investment, Pension, Annuities	0.10	0.30
Male Ombudsman	0.51	0.50
Contains Male Challenger	0.70	0.46
Contains Female Challenger	0.50	0.50

Notes: The means of “Contains Female Challenger” and “Contains Male Challenger” do not sum to one as challengers can be submitted by multiple individuals. Damages are in GBP in the data but for summary statistics are converted into 2026 USD. The Flesch variable represents the [Flesch \(1948\)](#) measure of readability

even split between male and female ombudsmen, 70% of cases involve a male challenger and 50% involve a female challenger (joint complaints are possible).

5.2.2 Analysis of the Damages Distribution

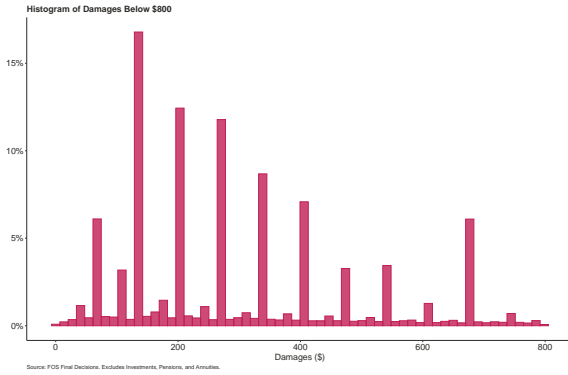
A central empirical feature of the FOS trial data is that the damage awarded (B in the model) has a long right tail. Appendix Figures A.2 and A.3 illustrate and describe tests rejecting log-normality and a simple (non-truncated) Pareto specification for the overall distribution. Instead, the distribution is characterized by the following two features.¹⁶

First, as shown in Figure 2a, the distribution has a pronounced discrete component at low values, with visible mass points at round numbers below approximately \$700. Around 68% of damages are below this cutoff. This is consistent with institutional practices that generate standardized awards for particular forms of harm (for example, modest amounts for distress and inconvenience) or rounding conventions in redress calculations.

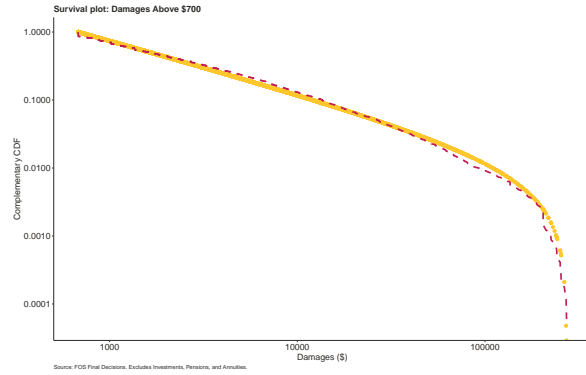
Second, above \$700, the distribution becomes continuous and strongly right-skewed, with rare but very large awards. In this region, as shown by Figure 2b, the empirical tail can be summarized by a truncated power-law type behavior up to around \$275,000 (the institutional cap is around

¹⁶In what follows I exclude pensions and investments as these typically involve larger damages and percentage terms, which I am not able to extract from the data.

FIGURE 2. Distribution of Damages in Upheld Cases



(A) Histogram of damages below \$800



(B) Survival plot of damages above \$700. Maroon dashed is data, solid orange dots is fitted model from truncated Pareto

\$580,000 but is very rarely used for non-pension cases).

The vast variation in damages implies that for the estimation of the structural model, the damages variable B must not be treated as a constant across cases. Further, because of the skewness in the distribution and its mixed nature, drawing from the non-parametric distribution is better than fitting a distribution to the data.

5.2.3 Heterogeneity in Uphold Rates, Damages, and Sentiment

Over Time Uphold rates vary over time. Uphold rates were around 23% in 2014 and then rose steadily to approximately 40% by 2024. This upward trend may reflect changes in product mix, complaint composition, ombudsman practice, or the nature of disputes reaching final decision, and it motivates robustness checks that estimate the model for different time periods. In real terms, mean damages fell between 2016 and 2020, before rising 33% between 2020 and 2025. As with upholding rates, these changes could reflect shifts in complaint composition, product mix, or the selection of disputes that reach final decision.

By Product Uphold rates vary meaningfully across product categories, from 29% in Banking to 40% in Consumer Credit. Conditional on a case being upheld, damages also vary sharply across products. Mean damages range from about \$1500 in Consumer Credit to roughly \$7000 in investments and pensions-related disputes. Because the pensions and investments product category differs from all other product categories, I exclude it when estimating the structural model. Average sentiment is positive towards consumers in investment/pensions cases, but negative in insurance cases. In general, sentiment and readability are relatively stable across products,

suggesting that this margin is not a significant source of heterogeneity.

While there are differences in average uphold rates and damages across product categories and years, ANOVAs of uphold rates and separately of damages reveal that the vast majority of variation in damages and uphold rates is within product-year pairs rather than across. So, even if I estimate the model separately by year and product category, I need sources of variation in cases that can explain differences in decisions and damages awarded across cases.

By Ombudsman and Challenger Gender Appendix Table A.1 shows the results of regressions of logged damages, average sentiment, and case decision on the gender of the ombudsman and challenger, along with fixed effects for the decision year and product category. Controlling for year and product category, there is no evidence of systematic differences in uphold rates by the gender of the ombudsman. However, damages are approximately 8% higher when the ombudsman is male than when the ombudsman is female. Also, male ombudsmen have a positive average sentiment, relative to a negative average sentiment for females, but decision letters written by male ombudsmen are slightly harder to read.

Uphold rates differ systematically by challenger gender composition: controlling for product category and year dummies, cases in which at least one challenger is male are approximately 8% (three percentage points) less likely to be upheld. However, conditional on being upheld, and controlling for product category and year dummies, cases that contain at least one male challenger have damages that are around 8% higher on average.¹⁷ Hence, while male challengers are less successful in their challenges, when they win, they are awarded more. Finally, controlling for product category and year dummies, the letter's tone is more positive when there is a male challenger, despite the lower uphold rate.

Heterogeneity Across Ombudsmen and Firms Beyond product and time variation, the data reveal substantial heterogeneity at the level of individual ombudsmen and respondent firms. Ombudsmen differ in their average uphold rates and in the typical damages they award in upheld cases, even within broad product categories. Figures A.4 and A.5 summarize heterogeneity across ombudsmen in uphold rates and damages.

Firms differ markedly in their uphold rates, which may reflect differences in underlying conduct, complaint composition, case handling, or selection into final decisions. Figures A.6 and A.7

¹⁷The relevant coefficients described here relating to the challenger gender are statistically significant no matter how I cluster standard errors, and coefficients on the gender of the ombudsman are significant apart from when I cluster by examiner.

summarize heterogeneity across firms in uphold rates and damages.

To provide a concrete benchmark, I also report total volumes in the sample, mean damages, and uphold rates for the fifteen most prevalent firms in the sample (Figures A.8, A.9, and A.10). There is wide variation in both uphold probabilities and conditional award sizes across the firms, and decisions are not concentrated at the top firms. The top fifteen firms are responsible for only 43% of cases, the HHI is 200, and hence the effective number of firms is 50.

6 Model Estimation and Parameter Estimates

In this section, I describe my estimation of the structural model as in Section 3.4. The model contains six variables to take to the data: damages B , lender fees f_L , high-merit and low-merit uphold probabilities e_G and e_I , the proportion of cases with $\theta = G$ given by π , and the cost of pursuit, K . The first three of the six are directly identified in the data. Regarding B , I observe its distribution for public cases. I assume this distribution is representative, and draw from the distribution of B in the data. The values of the lender fee f_L over time are also observed, ranging from approximately \$850 to \$1050 for the years I have data on. Regarding e_G , note that the equilibrium of the model I estimate only features trials when lenders are guilty. Hence, if the observed data matches the equilibrium of the model, the uphold rate is equal to the probability that a high-merit case is upheld. So e_G can be set externally as the uphold rate in the data (around 35%).

The remaining three parameters, $\psi = (e_I, \pi, K)$, have no observable counterparts and must be estimated inside the model. I estimate the model using simulated method of moments (SMM). For a given guess of ψ , I draw a set of values of B and θ , calculate equilibrium outcomes, and match three moments that are observed in the data. I draw on three moments on endogenous variables I can observe in the data. I describe these three moments $m_S = (m_1, m_2, m_3)$ in turn.

First, I match on the proportion of consumer issues that result in a customer complaint. I use figures from the *FCA Financial Lives Survey (FLS)*, which asks consumers directly whether they had an issue with a financial lender, and whether they complained to the lender about it.¹⁸ The decision of a consumer to complain about an issue depends directly on K but also on e_I and π , implying that this moment contains a signal on the parameters to estimate. The FLS contains data on this moment for the years 2017, 2022, and 2024.

¹⁸See [Financial Conduct Authority \(2025\)](#) for a technical description of the FLS and [Babina et al. \(2025\)](#) for an academic paper that uses the FLS data.

Second, I draw on FCA statistics on the average settlement amount in cases that are settled before a FOS case. The model implies that the settlement amount T is either equal to $e_I B$ or $e_G B$, with π and K governing (among other variables) the relative likelihood of these two offers. These FCA data are available each year from 2019 onwards.

Finally, the FOS publishes the number of overall FOS trials by quarter, which, together with FCA data on the number of complaints (which firms must log at the FCA by law), allows me to calculate the proportion of pursuits going to FOS trial as my third moment to match on. Aggregate data on the number of FOS trials are available from the 2022 calendar year onward. As a result, my three moments overlap for the years 2022 and 2024. I estimate the model separately for these two years and for 2022 and 2024 combined. The SSM estimator solves

$$\min_{\psi} (m(\psi) - m_S)' W (m(\psi) - m_S)$$

where $m(\psi)$ is the corresponding vector of moments to m_S from the simulated model when the parameter vector is ψ . For the weight matrix W , since I do not have microdata on the three moments, I use a diagonal matrix with weights equal to the inverse of the data moments squared, to place the three moments on the same order of magnitude.

6.1 Parameter Estimates

Now I describe the parameter estimates. I report the estimates pooling years 2022 and 2024 (parameter estimates for each individual year are similar). I estimate a low proportion of guilty cases in the overall population, with π estimated at 3.1%. This estimate implies that almost all cases are of low merit, aligning with the fact that most complaints are not escalated to FOS. Further, the generally low settlement amounts together with moderate damages imply that the FOS only upholds innocent cases very rarely, with e_I also estimated at 2.4%. Thirdly, the estimate of K is approximately £56.8.¹⁹ Finally, I fit the three moments well, with an SMM objective of 0.004 at the estimates, matching all three moments precisely.²⁰

¹⁹For comparison, this is equivalent to four hours of work for a worker on a minimum wage salary.

²⁰Because I do not have the underlying microdata that constructs the moments I use for estimation, I cannot calculate statistical standard errors. However, through simulating the model multiple times I can calculate the simulation error. The standard deviation of the estimates (e_I, π, K) over 50 simulations using different draws of fundamental shocks each time is (0.002, 0.006, 7.303). The estimates reported in this paragraph are means over 50 simulations.

7 Counterfactual Analysis

In the final part of the empirical analysis, I simulate the estimated model in the baseline regime and in several counterfactual scenarios, including changes to the fee structure, institution errors, and reducing frictions in making a challenge. While I analyze these counterfactuals in the FOS context, they correspond to the main design choices faced by redress institutions more broadly: (i) defendant-funded processes with full or partial refunds, (ii) no-fee processes with external funding, and (iii) refundable challenger deposits.

In each counterfactual I consider, I calculate the proportion of issues pursued by consumers, the proportion of trials conditional on pursuit, the total payments made by innocent and guilty lenders respectively, and the case costs and case revenues for the FOS institution. For FOS case revenues, I multiply the number of FOS cases by the lender case fee. For FOS case costs, I multiply the number of FOS cases by the FOS's reported marginal cost per case, given as \$1,400 for 2024/2025.

Table 2 reports the results. The values in the table correspond to one year's worth of potential consumer financial issues. In the baseline, approximately 25% of issues are pursued, and 94% of pursuits are settled before a FOS trial. Innocent lenders' payments to consumers and FOS fees amount to approximately \$1.3bn per year in total, whereas guilty lenders pay only \$947m. FOS case costs and revenues amount to around \$391m and \$247m, respectively. These figures are in line with those reported in the Ombudsman's budget document for 2024/25. As mentioned earlier, the gap between FOS case costs and revenues is covered by the annual levy on lenders.

In the first three counterfactuals I consider, I first remove errors ($e_I = 0, e_G = 1$), second turn off lender fees ($f_L = 0$), and third remove both lender fees and FOS errors. Removing FOS errors reduces costs for innocent lenders by 90%, as the low settlement cost now lies below the cost of pursuit, and they are never upheld in a trial. Guilty lenders' payments increase by more than double, as they cannot get away with a trial mistake, and the higher settlement transfer increases. Since trials increase, so do both FOS case costs and revenues by 58%.

Removing lender fees in trials in the presence of FOS errors reduces costs to both innocent and guilty lenders, as settlement becomes less likely and the cost of trials vanishes. As a result of lenders' increased bargaining power, there are 46% fewer pursuits in the first place. Interestingly, the decreased pursuits and increased trial probability (conditional on pursuit) cancel each other out, leaving the number of trials (and thus FOS case costs) unchanged, but, of course, the FOS case revenues become zero. Ultimately, removing FOS case fees alone would leave a hole in the FOS's budget and not improve efficiency.

However, removing FOS case fees alongside FOS errors would eliminate costs to innocent lenders and double payments from guilty lenders back to consumers. Yet the rise in trials against guilty lenders increases FOS costs, while the removal of lender fees eliminates FOS revenues. So, while this policy change can remove inefficiencies, it will not be budget-balancing and would require financing through higher levies or taxes. While innocent lenders would benefit from paying a higher levy to cover the loss in case revenues and rise in case costs, it is unclear how much investment (for example, in examiner training) would be required to remove FOS errors.

I also consider the effects of removing pursuit costs ($K = 0$). Naturally, in this case, all consumers pursue issues with lenders. The equilibrium is entirely in pure strategies as described in Section 3.4.1. In this case, the proportion of trials almost exactly matches the proportion of cases in the population with $\theta = G$, as π is below π^* and so settlement offers are almost entirely low. Hence, costs to guilty lenders increase by 30%. The increase in trials implies a 112% increase in FOS costs and FOS revenues. Hence, efforts to make pursuing the lender cheaper for consumers increase payments from lenders to consumers, but almost double the deficit in FOS costs relative to case revenues, and would require further funding that lenders would not be willing to cover.

The final counterfactual in Table 2 that I consider is the consumer deposit scheme, as explored theoretically in Section 4. I consider the extreme case in which lenders do not have to pay a case fee ($F_1 = 0$). As explained in the theoretical analysis, only high-merit cases will be pursued and all pursued high-merit cases will go to trial. The result is that innocent lenders save over \$1bn in transfers, but double guilty lenders' payments. FOS revenues fall to zero in this conservative case where lender fees are zero, and FOS costs over double, implying an increase in net operating costs of \$730m. Crucially, even in this conservative case where no revenues are raised from trials, the savings made by innocent lenders would more than cover the increased deficit coming from increased trials for guilty lenders. So, an increased levy on lenders in exchange for removing case fees would benefit innocent lenders and harm guilty lenders. In this extreme case, the model implies that the wish to keep the Ombudsman free of fees for consumers is creating a transfer of approximately \$1.3bn per year from innocent lenders to consumers with low-merit cases, and yet allowing guilty lenders to avoid nearly \$1bn per year in misconduct charges that a more efficient, but less affordable FOS design could achieve.

Finally, as is typical in counterfactual analysis from a structural model, I assume that all estimated parameters remain fixed, except those that are changed in the counterfactual itself. This assumption could be questioned in some of the counterfactual changes I consider, especially the rate of misconduct in the underlying population, π . For example, if lender fees f_L alone are removed, we may expect that rates of misconduct in the population would increase. Further, not only might the rates of misconduct change, but the size of the misconduct, B , may adapt as

TABLE 2. Counterfactual Results

Regime	Pursue Rate (%)	Trial Pursue (%)	I Cost	G Cost	FOS Cost	FOS Rev
Baseline	24.36	5.75	1309.72	947.12	391.90	247.56
No FOS Errors	11.41	19.65	133.91	2273.59	618.89	390.94
No Lender Fee ($f_L = 0$)	13.27	10.60	1136.69	699.56	391.90	0.00
No FOS Errors, $f_L = 0$	2.96	75.44	0.00	1882.65	618.89	0.00
Free Pursuit	100.00	2.96	1370.02	1232.14	829.53	524.00
Deposit Fee	3.13	100.00	0.00	1884.62	874.31	0.00

Notes: All costs and revenues are reported in millions of 2026 USD. No FOS errors means $e_I = 0$ and $e_G = 1$; Free pursuit means $K = 0$; Deposit Fee implements the counterfactual described in Section 4.3.

well. For example, if FOS errors are removed, we may expect that the distribution of B could decrease. Ultimately, these suggestions are speculative, and an encompassing model of lenders’ decisions to engage in misconduct (and the extent of misconduct) would be necessary to make more precise statements here.

8 Concluding Remarks

This paper develops and analyzes a model of the UK institution handling consumer financial redress for misconduct. The model includes selection into challenges by consumers, pre-trial settlement, adjudication errors, and procedural fees. I use the framework to clarify how fee design affects settlement incentives, selection into adjudication, and the efficiency of redress.

The analysis delivers three main insights. First, lender fees enable low-merit challengers to settle before an ombudsman decision. As a result, uphold rates in decided cases overstate population-level misconduct. Second, the fee structure generates an inherent tension between consumer affordability and efficiency, as keeping the process free for consumers creates incentives for opportunistic behavior by low-merit consumers. Third, introducing an upfront challenger fee for approaching the lender can screen out low-merit disputes, but it deters meritorious complaints and undermines the objective of free access for consumers. Counterfactual fee changes show that such policies would deliver savings of around \$1.3bn for non-liaible firms by reducing the volume of weak claims, but would increase the net costs of running the system by \$730m. If policy changes must be budget-balancing, the result implies that innocent lenders would benefit from a higher yearly levy in exchange for lower case fees.

Although the empirical setting is the UK financial ombudsman, the model’s economic insights and the affordability–efficiency tradeoff it formalizes apply to any redress institution whose operation is funded by one side of the disputes it adjudicates - a design feature shared by ombudsman schemes in many developed economies. As such, the insights of the model and the qualitative

conclusions from the counterfactuals extend to a broader set of institutions handling financial disputes.

There are two natural directions for future research. The first is to embed the redress system into a model of *ex ante* lender behavior, in which financial service providers endogenously choose whether to engage in misconduct in the first step, taking into account consumers' financial literacy, regulatory restrictions, and redress mechanisms. Such an extension would connect redress design directly to deterrence and allow welfare analysis that accounts for how institutional rules shape *ex ante* incentives. The second is to incorporate an explicit intermediary sector, modeling the role of professional representatives, who affect consumer participation, case selection, and settlement bargaining. Given that professional representatives were at the heart of the recent media and policy attention given to the Financial Ombudsman, understanding their role is a potentially important addition (Financial Ombudsman Service, 2025).

References

- AGARWAL, S., S. CHOMSISENGPHET, N. MAHONEY, AND J. STROEBEL (2014): "Regulating Consumer Financial Products: Evidence from Credit Cards," *The Quarterly Journal of Economics*, 130, 111–164, <https://doi.org/10.1093/qje/qju037>.
- ANNAN, F. (2025): "Misconduct and Reputation under Imperfect Information," *Journal of Political Economy*, 133, 1460–1496, <https://doi.org/10.1086/734133>.
- BABINA, T., S. BAHAJ, G. BUCHAK, F. DE MARCO, A. FOULIS, W. GORNALL, F. MAZZOLA, AND T. YU (2025): "Customer data access and fintech entry: Early evidence from open banking," *Journal of Financial Economics*, 169, 103950, <https://doi.org/10.1016/j.jfineco.2024.103950>.
- BANKS, J. S. AND J. SOBEL (1987): "Equilibrium Selection in Signaling Games," *Econometrica*, 55, 647–661, <https://doi.org/10.2307/1913604>.
- BEBCHUK, L. A. (1984): "Litigation and Settlement Under Imperfect Information," *The RAND Journal of Economics*, 15, 404–415, <https://doi.org/10.2307/2555448>.
- CHO, I.-K. AND D. M. KREPS (1987): "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 102, 179–221, <https://doi.org/10.2307/1885060>.
- DOU, Y., M. HUNG, G. SHE, AND L. L. WANG (2024): "Learning from peers: Evidence from disclosure of consumer complaints," *Journal of Accounting and Economics*, 77, <https://doi.org/10.1016/j.jacceco.2023.101620>.

- EGAN, M., G. MATVOS, AND A. SERU (2019): “The Market for Financial Adviser Misconduct,” *Journal of Political Economy*, 127, 233–295, <https://doi.org/10.1086/700735>.
- (2022): “When Harry Fired Sally: The Double Standard in Punishing Misconduct,” *Journal of Political Economy*, 130, 1184–1248, <https://doi.org/10.1086/718964>.
- (2024): “The Problem of Good Conduct among Financial Advisers,” *Journal of Economic Perspectives*, 38, 193—210, <https://doi.org/10.1257/jep.38.4.193>.
- (2025): “Arbitration with Uninformed Consumers,” *The Review of Economic Studies*, 92, 3888–3923, <https://doi.org/10.1093/restud/rdaf010>.
- ELIASON, P., R. LEAGUE, J. LEDER-LUIS, R. C. MCDEVITT, AND J. W. ROBERTS (2025): “Ambulance Taxis: The Impact of Regulation and Litigation on Health-Care Fraud,” *Journal of Political Economy*, 133, 1661–1702, <https://doi.org/10.1086/734134>.
- FCA (2020): “Payment protection insurance complaints deadline: Final Report,” *Financial Conduct Authority Research Publications*, <https://www.fca.org.uk/publication/ppi/payment-protection-insurance-complaints-deadline-final-report.pdf>.
- FCA AND FOS (2025): “CP25/22: Modernising the Redress System,” Consultation Paper CP25/22, Financial Conduct Authority and Financial Ombudsman Service, <https://www.fca.org.uk/publication/consultation/cp25-22.pdf>.
- FENN, P. AND N. RICKMAN (1999): “Delay and Settlement in Litigation,” *The Economic Journal*, 109, 476–491, <https://doi.org/10.1111/1468-0297.00458>.
- FINANCIAL CONDUCT AUTHORITY (2025): “Financial Lives 2024 Survey: Technical Report,” Technical report, Financial Conduct Authority, <https://www.fca.org.uk/publication/financial-lives/financial-lives-survey-2024-technical-report.pdf>.
- FINANCIAL OMBUDSMAN SERVICE (2025): “Financial Ombudsman Service to start charging professional representatives to refer cases,” Press release, <https://www.financial-ombudsman.org.uk/news/financial-ombudsman-service-start-charging-professional-representatives-refer-cases>.
- FLESCH, R. (1948): “A New Readability Yardstick,” *Journal of Applied Psychology*, 32, 221–233, <https://doi.org/10.1037/h0057532>.
- FREILICH, J., M. MEURER, M. SCHANKERMAN, AND F. SCHUETT (2024): “A New Approach to Patent Reform,” *UC Irvine Law Review*, 14, 351–403, <http://dx.doi.org/10.2139/ssrn.4357386>.

HOUSING OMBUDSMAN SERVICE (2024): “The Housing Ombudsman Scheme,” *Housing Ombudsman Service*, <https://www.housing-ombudsman.org.uk/app/uploads/2024/09/01.-Corporate-Housing-Ombudsman-Scheme.pdf>.

LEE, Y.-H. A. AND D. KLERMAN (2016): “The Priest-Klein hypotheses: Proofs and generality,” *International Review of Law and Economics*, 48, 59–76, <https://doi.org/10.1016/j.irl.2016.06.002>.

LOVE, B. J., S. P. MILLER, AND S. AMBWANI (2019): “Determinants of Patent Quality: Evidence from Inter Partes Review Proceedings,” *University of Colorado Law Review*, 90, 67, <http://dx.doi.org/10.2139/ssrn.3119871>.

LÜHRMANN, M., M. SERRA-GARCIA, AND J. WINTER (2015): “Teaching teenagers in finance: Does it work?” *Journal of Banking & Finance*, 54, 160–174, <https://doi.org/10.1016/j.jbankfin.2014.11.009>.

——— (2018): “The Impact of Financial Education on Adolescents’ Intertemporal Choices,” *American Economic Journal: Economic Policy*, 10, 309—32, <https://doi.org/10.1257/pol.20170012>.

LUSARDI, A. AND O. MITCHELL (2007): “Financial Literacy and Retirement Preparedness: Evidence and Implications for Financial Education,” *Business Economics*, 42, 35–44, <https://doi.org/10.2145/20070104>.

LUSARDI, A. AND O. S. MITCHELL (2008): “Planning and Financial Literacy: How Do Women Fare?” *American Economic Review: Papers and Proceedings*, 98, 413–17, <https://doi.org/10.1257/aer.98.2.413>.

MATCHAM, W. (2025): “Risk-Based Borrowing Limits and Interest Rates in Credit Card Markets,” *Unpublished Working Paper*, <http://dx.doi.org/10.2139/ssrn.4926974>.

MATCHAM, W. AND M. SCHANKERMAN (2025): “Screening Property Rights for Innovation,” <http://dx.doi.org/10.2139/ssrn.4519999>.

NELSON, S. T. (2025): “Private Information and Price Regulation in the US Credit Card Market,” *Econometrica*, 93, 1371–1410, <https://doi.org/10.3982/ECTA18063>.

PRIEST, G. L. AND B. KLEIN (1984): “The Selection of Disputes for Litigation,” *The Journal of Legal Studies*, 13, 1–56, <http://dx.doi.org/10.1086/467732>.

RINKER, T. W. (2021): *sentimentr: Calculate Text Polarity Sentiment*, <https://doi.org/0.32614/CRAN.package.sentimentr>.

- RUTLEDGE, S. L. (2010): “Consumer Protection and Financial Literacy: Lessons from Nine Country Studies,” Policy Research Working Paper 5326, The World Bank, Washington, DC.
- SALANT, S. W. (1984): “Litigation of Settlement Demands Questioned by Bayesian Defendants,” Working Papers 516, California Institute of Technology, Division of the Humanities and Social Sciences.
- SCHANKERMAN, M. AND F. SCHUETT (2022): “Patent Screening, Innovation, and Welfare,” *The Review of Economic Studies*, 89, 2101–2148, <https://doi.org/10.1093/restud/rdab073>.
- SOBEL, J. (1989): “An Analysis of Discovery Rules,” *Law and Contemporary Problems*, 52, 133–159, <https://doi.org/10.2307/1191900>.
- SPENCE, M. (1973): “Job Market Signaling,” *The Quarterly Journal of Economics*, 87, 355–374, <https://doi.org/10.2307/1882010>.
- SPIER, K. (1997): “A Note on the Divergence between the Private and the Social Motive to Settle under a Negligence Rule,” *The Journal of Legal Studies*, 26, 613–621, <https://doi.org/10.1086/468009>.
- SPIER, K. E. (1992a): “The Dynamics of Pretrial Negotiation,” *The Review of Economic Studies*, 59, 93–108, <https://doi.org/10.2307/2297927>.
- (1992b): “Incomplete Contracts and Signalling,” *The RAND Journal of Economics*, 23, 432–443, <https://doi.org/10.2307/2555872>.
- (1994): “Pretrial Bargaining and the Design of Fee-Shifting Rules,” *The RAND Journal of Economics*, 25, 197–214, <https://doi.org/10.2307/2555827>.
- UNITED STATES V. BLUE BELL CREAMERIES L.P. (2020): “United States District Court,” W.D. Texas, , No. 5:20-cr-00370.
- UNITED STATES V. GOOGLE LLC AND YOUTUBE, LLC (2019): “United States District Court,” D.D.C., No. 1:19-cv-02642.
- UNITED STATES V. HOLMES (2018): “United States District Court,” N.D. California, No. 5:18-cr-00258.
- UNITED STATES V. VOLKSWAGEN AG (2017): “United States District Court,” E.D. Michigan, No. 3:16-cr-00203.
- VAISHNAV, D., M. NEETHINAYAGAM, A. KHAIRE, AND J. WOO (2024): “Predictive Analysis of CFPB Consumer Complaints Using Machine Learning,” <https://doi.org/10.48550/arXiv.2407.06399>.

A Additional Tables and Figures

TABLE A.1. Regression Results

Variable	(1) Uphold	(2) Log(Damages)	(3) Sentiment
MALE OMBUDSMAN	0.008 (0.003) [0.008]	0.082 (0.020) [0.060]	0.002 (0.000) [0.002]
A MALE CHALLENGER	-0.025 (0.007) [0.007]	0.077 (0.021) [0.022]	0.016 (0.001) [0.001]
Decision Year Dummies	Yes	Yes	Yes
Product Dummies	Yes	Yes	Yes
N	350,158	54,937	350,170

Notes: Column (1) shows estimates from a regression of a binary variable equal to one if the case is upheld, against dummies for decision year, product group, and dummies for the gender of the ombudsman and challenger. In column (2), the dependent variable is the log of the damages; the regressors are the same as column (1). In column (3), the dependent variable is the average sentiment of the decision letter; the regressors are the same as columns (1) and (2). Standard errors in curly parentheses are clustered at the firm level and standard errors in square parentheses are clustered at the examiner and firm level.

FIGURE A.1. Constructing case-level data from the FOS public decisions portal and PDF decision letters.

The screenshot shows a public decision entry with the following elements: a title 'Decision Reference DRN-5425291' where 'DRN-5425291' is highlighted in an orange box; a date '7 Apr 2025' in a red box; a firm name 'American Express Services Europe Limited' in a blue box; an outcome 'Upheld' in a green box; and a product category 'Banking and Payments' in a purple box. Below this is a summary of the complaint: 'DRN-5425291 The complaint Mr N complains that American Express Services Europe Limited (Amex) reduced his credit limit and that resulted in him going over the limit. He also says they wrongly reported his account status to the Credit Reference Agencies (CRA's). What happened Mr N has a ... (2 pages)'. At the bottom is a red link labeled 'View decision'.

(A) Public FOS decisions portal: ID, date, firm, outcome, and product category.

My final decision

For the reasons I've given above, I uphold this complaint in part and tell American Express Services Europe Limited to pay Mr N £100 in compensation for the distress and inconvenience caused.

Under the rules of the Financial Ombudsman Service, I'm required to ask Mr N to accept or reject my decision before 5 May 2025.

Phillip McMahon
Ombudsman

(B) Extracting consumer honorifics, damages, and ombudsman identity from the decision text.

FIGURE A.2. Damage QQ Plot: Log-Normal

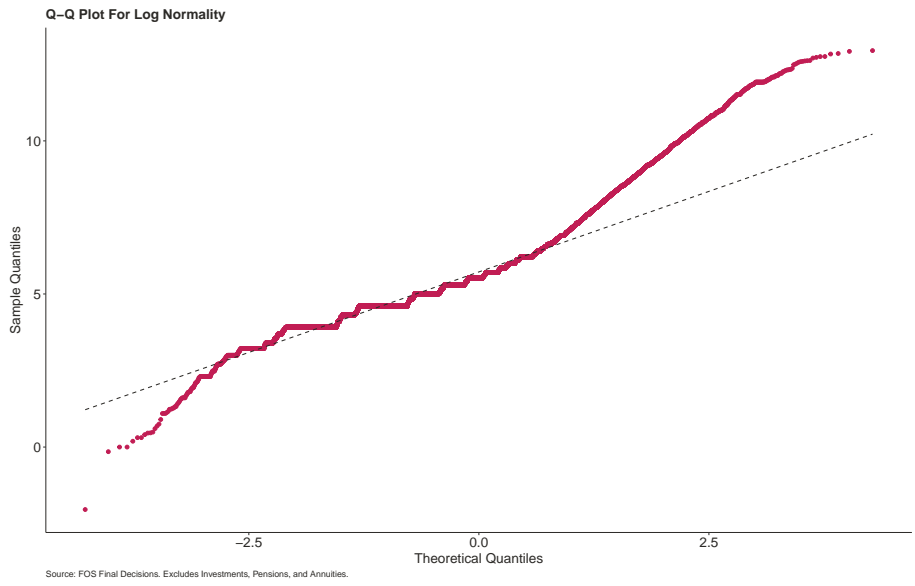


FIGURE A.3. Damage QQ Plot: Pareto




FIGURE A.4. Histogram of Ombudsmen Uphold Rates

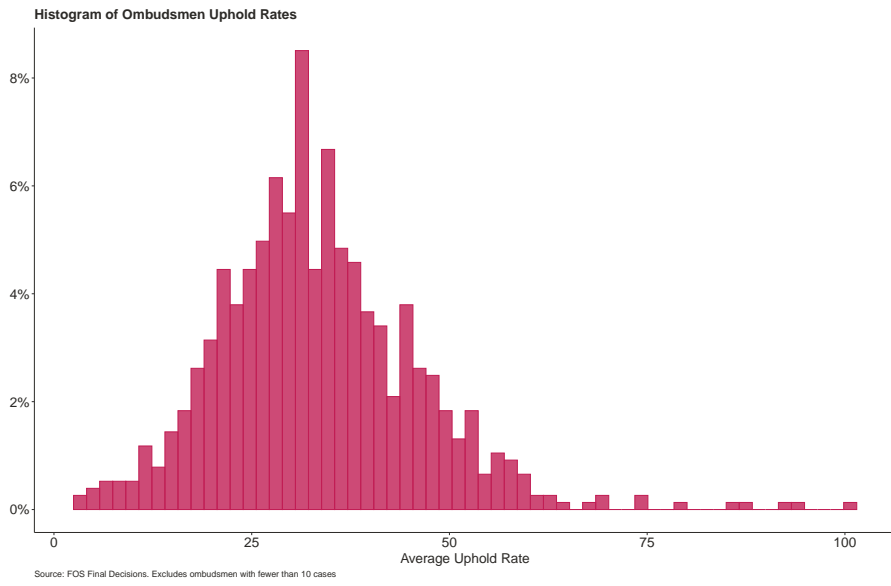


FIGURE A.5. Histogram of Ombudsmen Mean Log Damages




FIGURE A.6. Histogram of Firm Uphold Rates




FIGURE A.7. Histogram of Firm Mean Log Damages




FIGURE A.8. Histogram of Firm Total Volumes: Top 15 Firms




FIGURE A.9. Histogram of Firm Uphold Rates: Top 15 Firms





FIGURE A.10. Histogram of Firm Damages: Top 15 Firms



B Additional Theoretical Analysis for the Model

I consider several extensions, so it helps to summarize the subsections that follow, as in Table B.2. In Section 3.2 in the text, I consider the case where challenger and lender observe merit. Section B.1 analyzes the case in which both challenger and lender do not observe merit. Sections B.2.1 and B.2.2 cover the case of the challenger making an offer, considering the separating and pooling equilibria respectively. Finally, B.3 covers the case of the challenger not knowing the merit. Section B.3.1 analyzes this environment with the challenger making an offer, and Sections B.3.2 and B.3.3 the case of the lender making the offer, again covering separating and pooling equilibria respectively.

B.1 Model Equilibrium: Challenger and Lender Don't Observe Merit

When both challenger and lender do not observe merit, the situation is similar to the complete information case except that both must work with their expected likelihood of winning a trial from a trial, rather than their actual likelihood. As such, the lender stands to lose $f_L + \bar{e}B$ from a trial and the challenger stands to gain $\bar{e}B$, where $\bar{e} = \pi e_G + (1 - \pi)e_I$. Thus the settlement region is $[\bar{e}B, f_L + \bar{e}B]$, and the challenger pursues if $\bar{e}B > K$. As is clearly seen, this case only differs from the complete information case in using \bar{e} rather than e_θ , though because the challenger does not know their type, as was discussed in Section 3.5, it is not possible to deter low-merit types alone from challenging.

TABLE B.2. Summary of extensions

One-line summary of the case	Section
Challenger and lender know merit	3.2
Challenger and lender do not know merit	B.1
Challenger knows merit, lender does not	
Lender makes settlement offer & $K < e_I B$	3.4.1
Lender makes settlement offer & $K \in [e_I B, e_G B]$	3.4.2
Challenger makes offer, separating equilibrium	B.2.1
Challenger makes offer, pooling equilibrium	B.2.2
Lender knows merit, challenger does not	
Challenger makes offer	B.3.1
Lender makes offer, separating equilibrium	B.3.2
Lender makes offer, pooling equilibrium	B.3.3

Notes:

B.2 Model Equilibrium: Lender Doesn't Observe Merit

B.2.1 Challenger Offers Settlement: Separating Equilibrium

An equilibrium with the challenger proposing the settlement level involves decisions to challenge, settlement offers T_I, T_G , lender posteriors of $\theta = G$ updated based on T , denoted $\mu(T)$, and a rejection rule $r(T) \in \{\text{Accept}, \text{Reject}\}$. In a separating equilibrium, $T_I \neq T_G$. I start with Perfect Bayesian Equilibrium (PBE), which requires that both types of challengers do not want to deviate to any other T , $r(T)$ is optimal for all T given μ , and that beliefs are updated using Bayes's rule on equilibrium path, that is, $\mu(T_I) = 0 = 1 - \mu(T_G)$. Result 1 that follows characterizes the separating equilibrium I focus on.

Result 1. *The following strategy is a Perfect Bayesian Equilibrium if $K < f_L + e_I B$ and $f_L < (e_G - e_I)B$:*

- *Both types pursue the lender*
- $T_I = f_L + e_I B$
- $T_G = f_L + e_G B$
- $r(T) = \text{Reject}$ if and only if $T > T_I$

- $\mu(T) = 0$ for $T < T_G$ and $\mu(T) = 1$ for $T \geq T_G$.²¹

The following proves Result 1:

- For a challenger of type I , any deviation to $T < T_I$ is also accepted but yields a lower payoff. Any deviation to $T > T_I$ is rejected and thus yields a payoff $e_I B$ which is lower than $f_L + e_I B$.
- For a challenger of type G , any deviation to $T > T_I$ is also rejected, so their payoff is not improved. Any deviation $T \leq T_I$ is accepted so the only deviation to consider is to the highest of that range, which is to T_I . This is not optimal provided $f_L + e_I B < e_G B$ which is equivalent to $f_L < (e_G - e_I)B$, as I assume.
- The lender optimally accepts $T \leq T_I$ because rejecting yields cost $f_L + e_I B \geq T_I$
- For $T \neq T_G$, the lender optimally rejects $T > T_I$ because in this range they believe they are innocent with probability 1, so they would rather pay $f_L + e_I B$ than $T > f_L + e_I B$
- The lender is indifferent between accepting and rejecting T_G , because they believe they are guilty with probability 1 when the challenger offers T_G , so will pay $f_L + e_G B$ whether they reject or accept.²²
- Both types pursue the lender by assumption as $K < f_L + e_I B < e_G B$.

As is quite clear from the form of the equilibrium, there is a continuum of alternative separating PBEs differing in T_G and μ . Finally, while this equilibrium is not ruled out by the intuitive criterion of [Cho and Kreps \(1987\)](#), it cannot survive the D1 Divine Equilibrium of [Banks and Sobel \(1987\)](#).²³

B.2.2 Challenger Offers Settlement: Pooling Equilibrium

A pooling equilibrium involves $T_I = T_G = T_C^{\text{pool}}$ and lender posterior beliefs of $\theta = G$ satisfying $\mu(T_C^{\text{pool}}) = \pi$. I consider the pooling equilibrium that survives refinements. The pooling equilibrium is characterized by the following conditions:

Result 2. *Let $\bar{T}_C = f_L + \pi e_G B + (1 - \pi) e_I B$. The following is a pooling equilibrium if $K < \bar{T}_C$ and $f_L \geq (1 - \pi)(e_G - e_I)B$:*

²¹An alternative for a PBE that doesn't survive the Intuitive Criterion has $\mu(T) = 0$ for $T > T_G$.

²²This is a tie-break among the best responses that is required for the low-merit type's incentive constraint.

²³I am not able to find a separating equilibrium surviving D1 refinement.

- $T_C^{pool} = \bar{T}_C$
- $\mu(T_C^{pool}) = \pi$, $\mu(T) = 0$ for all $T \neq T_C^{pool}$
- $r(T_C^{pool}) = \text{Accept}$; for all $T \neq T_C^{pool}$, $r(T) = \text{Accept}$ if and only if $T \leq f_L + e_I B$.

For $T \neq \bar{T}_C$, the lender optimally rejects $T > f_L + e_I B$ because for these T , they believe they are innocent with probability 1, so trial gives a cost of $f_L + e_I B < T$. The lender optimally accepts $T < f_L + e_I B$ for the same logic. The lender is indifferent between accepting and rejecting T_C^{pool} .

The G-type challenger does not want to deviate to $T > f_L + e_I B$ as that involves a trial which gives payoff $e_G B \leq \bar{T}_C$ by assumption. The G-type challenger does not want to deviate to $T < f_L + e_I B$ as this is accepted and lower than T_C^{pool} .²⁴ The same is true for an I-type challenger. The I-type challenger does not want to deviate to $T > f_L + e_I B$ as that involves a trial which gives payoff $e_I B < e_G B < T_C^{pool}$.

The equilibrium requires $f_L > (1 - \pi)(e_G - e_I)B$: the intuition is that if the lender's FOS costs are sufficiently large, the G-type challenger prefers to settle and extract some of those FOS fees and be pooled with innocent types than to separate themselves from an I-type and obtain the trial payoff.

B.3 Model Equilibrium: Challenger Doesn't Observe Merit

As I explain in text, I prefer the model with the challenger observing the merit and the lender not, as it is not possible to work through policies to deter low-merit challenges when challengers do not know their merit. Nevertheless, I solve the model in this case to show that the types of equilibria in the cases with the challenger not observing the merit are the same as those described in text.

B.3.1 Challenger Offers Settlement

For simplicity, assume $K < e_I B$ so that for all π , the challenger wants to pursue. For the same reasons in text, the challenger will consider two possible offers: $T = f_L + e_I B$, in which case there will be settlement by both types, or $T = f_L + e_G B$, in which case only type G will accept. In the former offer, the expected payoff for the challenger is $f_L + e_I B$. In the latter offer, the expected payoff is

$$\pi(f_L + e_G B) + (1 - \pi)e_I B$$

²⁴This result follows from the fact that $e_G B < f_L + \pi e_G B + (1 - \pi)e_I B \iff f_L > (1 - \pi)(e_G - e_I)B$

After some algebra, it becomes clear that the both types accept the settlement offer of $f_L + e_I B$ if and only if

$$\pi < \frac{f_L}{\Delta_e B + f_L}$$

Since f_L is larger than B in most cases, this is the empirically likely outcome, and so settlement with G types only is empirically unlikely in this case.

B.3.2 Lender Offers Settlement: Separating Equilibrium

An equilibrium with the lender proposing the settlement level involves decisions to challenge, settlement offers T_I, T_G , challenger posteriors of $\theta = G$ updated based on T , denoted $\mu(T)$, and a rejection rule $r(T) \in \{\text{Accept}, \text{Reject}\}$. In a separating equilibrium, $T_I \neq T_G$. Result 3 that follows characterizes the separating Perfect Bayesian Equilibrium I focus on.

Result 3. *The following strategy is a Perfect Bayesian Equilibrium if $K < e_I B$ and $f_L < (e_G - e_I)B$:*

- *Both types pursue the lender*
- $T_I = 0$
- $T_G = e_G B$
- $r(T) = \text{Accept}$ if and only if $T \geq e_G B$
- $\mu(0) = 0$ and $\mu(T) = 1$ for $T > 0$.

The key part for the proof is to ensure that the innocent lender prefers trial than mimicking the guilty and paying the high settlement. This requires that the trial payoff is less than the settlement amount, i.e.

$$f_L + e_I B < e_G B \iff f_L < \Delta_e B$$

However, it is worth noting that (i) at current trial fees, for most cases in the data, this condition will not hold, and (ii) more generally that this equilibrium relies on extreme off-path beliefs, namely that *any* positive offer implies a guilty type with probability 1.

B.3.3 Lender Offers Settlement: Pooling Equilibrium

A pooling equilibrium involves settlement offers $T_I = T_G = T_L^{\text{pool}}$ and challenger posterior beliefs of $\theta = G$ satisfying $\mu(T_L^{\text{pool}}) = \pi$. A pooling equilibrium is characterized by the following conditions:

Result 4. Let $\bar{T}_L = \pi e_G B + (1 - \pi) e_I B = \bar{e} B$. The following is a pooling equilibrium if $K < \bar{T}_L$ and $f_L > \pi \Delta_e B$:

- $T_L^{pool} = \bar{T}_L$
- $\mu(T) = 1$ for all $T < T_L^{pool}$, $\mu(T_L^{pool}) = \pi$, $\mu(T) = 0$ for all $T > T_L^{pool}$
- Rejection rule: $r(T) = \text{Accept}$ if and only if $T \geq T_L^{pool}$

The key condition to check is whether the innocent lender wants to deviate to a lower settlement fee. Deviating gives a cost of $f_L + e_I B$, which exceeds \bar{T}_L if $f_L > \pi \Delta_e B$, as the result stipulates.

C Formal Analysis of Model Reforms

C.1 2025 Fee Reform Analysis

This subsection formalizes the claim in Section 4 that the 2025 FOS fee changes cannot screen out low-merit complaints. Under the 2025 fee regime, challengers face a case fee f_C , which is partially refunded when the complaint is upheld. I parameterize this by $\gamma_C \in [0, 1]$: the challenger pays f_C if the complaint is not upheld, and pays $(1 - \gamma_C)f_C$ if the complaint is upheld. Thus, the expected fee paid by the challenger at trial is $f_C(1 - e\gamma_C)$.

Similarly, the lender faces a case fee f_L and receives a partial refund when the complaint is *not* upheld. With refund parameter $\gamma_L \in [0, 1]$, the lender pays f_L if the complaint is upheld and pays $(1 - \gamma_L)f_L$ if it is not upheld. The expected lender fee paid at trial is therefore $f_L \Gamma_L(e)$ where $\Gamma_L(e) = (1 - \gamma_L) + \gamma_L e$.

Challenger’s trial value. Given success probability e , the challenger’s expected payoff from trial is

$$V_C(e) = eB - f_C(1 - e\gamma_C).$$

Thus the challenger’s minimum acceptable transfer (the challenger’s “reservation settlement”) is

$$\bar{T}(e) = V_C(e) = eB - f_C(1 - e\gamma_C). \tag{C.1}$$

In particular, for high-merit disputes,

$$\bar{T}_G \equiv \bar{T}(e_G) = e_G B - f_C(1 - e_G \gamma_C). \tag{C.2}$$

Lender’s preference for settlement. The lender’s expected cost from trial (given success probability e) is the expected damages plus the expected case fee:

$$V_L(e) = eB + f_L \Gamma_L(e). \tag{C.3}$$

Since $f_L \Gamma_L(e) > 0$ and $f_C(1 - e\gamma_C) \geq 0$, comparing (C.3) and (C.1) gives

$$V_L(e) - \bar{T}(e) = f_L \Gamma_L(e) + f_C(1 - e\gamma_C) > 0. \quad (\text{C.4})$$

Therefore, for any $e \in (0, 1)$ the lender strictly prefers paying the challenger's reservation settlement $\bar{T}(e)$ to going to trial.

Result 5 (No screening under the 2025 fee changes). *Under the 2025 fee changes, there is no equilibrium in which high-merit challengers pursue while low-merit challengers do not.*

Proof. Suppose a high-merit challenger has initiated bargaining with the lender. If the lender offers $T \geq \bar{T}_G$, the high-merit challenger prefers settlement to trial by definition of (C.2). By (C.4), the lender strictly prefers paying $T = \bar{T}_G$ to facing trial. Hence, whenever a high-merit challenger is present, the lender has a strict incentive to make a settlement offer that is accepted by the high-merit challenger; in particular, trial cannot be sustained as an equilibrium outcome for high-merit challengers.

Now consider screening: to screen low-merit challengers, the equilibrium would need to make imitation unattractive, i.e., it would need to induce behavior by high-merit challengers that low-merit challengers would not mimic. But since high-merit challengers always settle in equilibrium, and since low-merit challengers have a weakly lower reservation settlement (because $e_I < e_G$ implies $\bar{T}(e_I) \leq \bar{T}(e_G)$ from (C.1)), any settlement that is acceptable to high-merit challengers is also acceptable to low-merit challengers. Thus settlement behavior cannot separate types.

Therefore, the 2025 fees cannot generate screening by inducing only high-merit challengers to proceed while deterring low-merit challengers. Instead, the effect of increasing f_C and introducing refunds operates through (C.1): it lowers challengers' outside option from trial, slackens the settlement condition, and induces more disputes to settle prior to final adjudication. \square

C.2 Challenger Deposit Fee

This subsection proves that a *refundable challenger deposit* can generate screening, at the cost of violating the “free at point of use” principle.

Deposit scheme. Suppose that to pursue a complaint to a final FOS decision, the challenger must pay an upfront deposit $F_2 > 0$. The deposit is refunded if the complaint is upheld and is forfeited otherwise. The lender faces a fee $F_1 \geq 0$ that is paid when the complaint is upheld.

Trial values under the deposit scheme. If the challenger proceeds to trial, the expected payoff is

$$V_C^{\text{trial}}(e) = eB - (1 - e)F_2 - K,$$

since the deposit is refunded with probability e and forfeited with probability $1 - e$. If instead the parties settle for transfer T , and the deposit is forfeited, the challenger receives

$$V_C^{\text{settle}} = T - F_2 - K.$$

Forcing trial for high-merit complaints. To sustain an equilibrium in which high-merit complaints proceed to trial (so that low-merit complaints do not find imitation attractive), it must be that a high-merit challenger weakly prefers trial to settlement:

$$T - F_2 \leq e_G B - (1 - e_G)F_2 \iff T \leq e_G(B + F_2) =: \check{T}$$

The lender's expected cost of trial in a high-merit dispute is

$$V_L^{\text{trial}}(e_G) = e_G(B + F_1),$$

because the lender pays damages B and fee F_1 only when the complaint is upheld. If the lender were to settle, it must offer at least \check{T} to prevent the high-merit challenger from going to trial. Hence the lender prefers trial to settlement when $\theta = G$ if

$$e_G(B + F_1) \leq e_G(B + F_2) \iff F_1 \leq F_2. \quad (\text{C.5})$$

Participation and screening. For a high-merit challenger to pursue the complaint at all, the expected trial value must be nonnegative:

$$e_G B - (1 - e_G)F_2 - K > 0 \iff F_2 < \frac{e_G B - K}{1 - e_G}. \quad (\text{C.6})$$

For a low-merit challenger *not* to pursue, the expected trial value must be nonpositive:

$$e_I B - (1 - e_I)F_2 - K \leq 0 \iff F_2 \geq \frac{e_I B - K}{1 - e_I}. \quad (\text{C.7})$$

Combining (C.6) and (C.7) yields an interval of deposits that separates types:

$$F_2 \in \left[\frac{e_I B - K}{1 - e_I}, \frac{e_G B - K}{1 - e_G} \right). \quad (\text{C.8})$$

This interval is nonempty whenever $K < B$ (and $e_G > e_I$).

Result 6 (Screening with a refundable deposit). *Suppose $e_G > e_I$ and $K < B$. If the refundable deposit F_2 satisfies (C.8) and the lender fee satisfies $F_1 \leq F_2$, then there exists an equilibrium in which high-merit challengers pursue and proceed to trial, while low-merit challengers do not pursue.*

Proof. Let F_2 satisfy (C.8). Then (C.6) holds, so high-merit challengers strictly prefer to pursue. At the same time, (C.7) holds, so low-merit challengers weakly prefer not to pursue.

Given $F_1 \leq F_2$, the lender prefers trial to any settlement offer that would keep a high-merit challenger from going to trial, by (C.5). Hence, the lender does not offer such settlements, and high-merit challengers proceed to trial. Since low-merit challengers have nonpositive expected value from trial (and cannot profitably obtain settlement without pursuing), they do not pursue. This yields the separating outcome stated in the lemma. \square