

NBER WORKING PAPER SERIES

COPYRIGHT ENFORCEMENT:
EVIDENCE FROM TWO FIELD EXPERIMENTS

Hong Luo
Julie Holland Mortimer

Working Paper 22082
<http://www.nber.org/papers/w22082>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2016

The authors thank James Dana, Alberto Galasso, Shane Greenstein, Michael Grubb, William Hickman, Gastón Llanes, Mike Luca, Cynthia Montgomery, Felix Oberholzer-Gee, Jason Snyder, Scott Stern, Ricardo Perez-Truglia, Dennis Yao and seminar participants at the Canadian Conference on the Economics of Innovation and Entrepreneurship, Harvard Business School, MIT Sloan, Microsoft Research, NBER Economics of Digitization Conference, NBER Productivity Lunch, Stern School of Business, the Strategic Research Forum, and Washington University in St. Louis for helpful comments. We thank the Agency for providing the proprietary data and for their collaboration on the experiments. We also thank Danielle Wedde and Esther Yan for their excellent research assistance. All errors and omissions are ours. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Hong Luo and Julie Holland Mortimer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Copyright Enforcement: Evidence from Two Field Experiments
Hong Luo and Julie Holland Mortimer
NBER Working Paper No. 22082
March 2016, Revised August 2016
JEL No. L0,M2,O30,O34

ABSTRACT

Effective dispute resolution is important for reducing private and social costs. We study how resolution responds to changes in price and communication using a new, extensive dataset of copyright infringement incidences by firms. The data cover two field experiments run by a large stock-photography agency. We find that substantially reducing the requested amount generates a small increase in the settlement rate. However, for the same reduced request, a message informing infringers of the price reduction and acknowledging the possible unintentionality generates a large increase in the settlement rate; including a deadline further increases the response. The small price effect, compared to the large message effect, can be explained by two countervailing effects of a lower price: an inducement to settle early, but a lower threat of escalation. Furthermore, acknowledging possible unintentionality may encourage settlement due to the typically inadvertent nature of these incidences. The resulting higher settlement rate prevents additional legal action and significantly reduces social costs.

Hong Luo
Harvard University
Morgan Hall 241
Soldiers Field Road
Boston, MA 02163
hluo@hbs.edu

Julie Holland Mortimer
Department of Economics
Boston College
140 Commonwealth Avenue
Chestnut Hill, MA 02467
and NBER
julie.mortimer.2@bc.edu

1 Introduction

Disputes involving economic damages, such as breach of contract, overdue payment, and violation of property rights, are commonplace in the economy. Effectively resolving these disputes can save time and stress, and reduce legal expenses. The threat of legal sanction may be credible only in some circumstances (e.g., it is impractical when monetary claims are small), and it may not always be desirable even if effective, because a litigious reputation may drive away customers. Encouraging collaborative outcomes in these situations may require alternative, creative methods.

Copyright enforcement is an interesting example of this problem and is important in the management of intellectual property (IP). While digital technologies have facilitated piracy in recent years, they have also substantially lowered the costs of detecting infringement, making enforcement an increasingly relevant option.¹ How to do so effectively, however, remains a challenge. Monetary amounts are generally small, and infringers are numerous and dispersed. Many content producers want to avoid a litigious reputation because it may hurt their primary source of revenue from the legitimate market. Legal actions against file-sharing networks and infringing individuals, notably by trade associations in the music and movie industries, have been controversial and have met with limited success. In this paper, we study the effectiveness of new approaches to enforcement on copyright infringement settlement outcomes. In particular, we focus on two aspects of enforcement methods: (1) the requested settlement amount (pricing), and (2) communication.

Our empirical application comes from the stock-photography industry.² Stock photos are pre-shot images that are readily available for licensing. On behalf of photographers, agencies such as Corbis Images and Getty Images manage, market, and license images to business customers. An interesting feature that distinguishes our context from the prior piracy literature is the prevalence of uninformed infringement (Luo and Mortimer (2016)). As digital images are easily accessible online through search engines and social media sites, users often have limited knowledge about the legal obligation of using an image. Furthermore, the frequent use of third-party advertisers and web designers results in many indirect instances of infringement, for which end-users are ultimately liable. We use data provided by one of the leading agencies in the industry (hereafter, the Agency), which monitors the online use of a small set of its most expensive images and pursues settlement for instances of infringement carried out for commercial purposes. The dataset contains detailed information on all infringement instances for which the Agency pursued settlement in the U.S. and Canada between October 2013 and March 2014, including the requested amount, the characteristics of

¹Some recent examples of anti-piracy services that have been developed by private companies include Digimarc in the publishing industry and MarkMonitor (purchased by Thomson Reuters in 2012) in industries such as movies, music, software, and publishing.

²Glückler and Panitz (2013) estimates that the global revenue of the stock-photography industry was \$2.88 billion in 2011. In comparison, the revenues of ASCAP and BMI, the largest two performance rights organizations that collect royalties on behalf of copyright holders in the music industry for public use of their works, were \$1B each in 2012.

infringing firms, and settlement outcomes.

Our data are novel for three reasons. First, previous empirical studies on settlement have depended primarily on court filings, which capture only a small percentage of disputes.³ Our sample is ‘unselected’ in that it includes all instances of infringement for which the Agency pursues settlement, and the data capture the earliest stage of the enforcement process. This is important because the set of infringers (and their settlement behavior) at this stage are likely to differ from those in the cases that reach litigation. Second, it is rare for field data to contain exogenous variation in enforcement methods, and laboratory experiments may be hard to generalize to real-world settings. Our data are generated by two field experiments, and include exogenous variation in both pricing and communication. Third, infringing firms in our sample are largely representative of the U.S. firm population based on observable characteristics.⁴ This is valuable because it documents the prevalence of infringement among businesses and provides insight into the resolution of disputes between firms, rather than between individuals, as typically studied in lab experiments.

Prior to either of the two experiments, the Agency’s settlement offer is the sum of two components: a baseline licensing fee for each image and an additional \$400 per-image amount to recover the Agency’s enforcement expenses. The baseline licensing fee depends on the infringing use (e.g., a high-resolution image posted on a website’s home page) and is a weighted average of use-related list prices on the Agency’s licensing site. The two price components are not broken down separately in the Agency’s letters to infringers. In the *pricing* experiment, the Agency exogenously varies the requested settlement amount by eliminating the \$400 per-image enforcement cost in some cases but not in others. The wording of the letter is otherwise kept identical, and no information on the price reduction is provided. Baseline licensing fees range from \$380 to \$825 per image, so removing the enforcement cost is economically significant.

After the pricing experiment, the Agency removes the enforcement cost for all cases. After this policy change, we worked with the Agency on the second experiment—the *message* experiment. In this experiment, an extra statement is inserted in the letters sent to infringers in four treatment groups, with no extra statement added for the control group. The first two messages explain that the Agency has forgiven a \$400 per-image amount, either as a waiver of enforcement costs or as a discount. To justify the reduction, the messages acknowledge that the infringement may have been unintentional. The last two treatment messages are the same as the first two, but with a deadline after which the forgiven amount is added back. We do not have a treatment message that explicitly increases the threat of legal sanction, because the Agency moved

³Trubek et al. (1983) estimate that 90 percent of civil disputes are settled without filing suit. In our setting, according to Lex Machina (a legal analytics database), the Agency filed 12 cases in the U.S. courts between 2008 and 2013. In contrast, the Agency pursues settlement for about 1,000 cases per month in the U.S., implying a litigation probability of about 0.02 percent.

⁴According to the 2012 Statistics of U.S. Businesses, 62 percent of U.S. firms have 0-4 employees, 17 percent have 5-9, 10 percent have 10-19, 8.6 percent have 20-99, 1.4 percent have 100-499, and 0.3 percent have more than 500. The corresponding percentages in our data are 50, 17.6, 12, 15, 5.1, and 0.89.

away from any enforcement approach that was explicitly threatening.

The pricing experiment shows that after 30 days, removing the \$400 per-image charge increases the settlement probability by three percentage points (from 12 to 15 percent, marginally significant) for an average case. However, this increase is offset by the significantly lower revenue for those cases that do settle. As a result, the expected revenue per case is slightly lower for the lower-priced group (although the difference is not statistically significant). In contrast, the message experiment, after 30 days, shows that adding either of the first two messages increases settlement probability by 12 percentage points and expected revenue by \$65 per case, both representing an increase of over 80 percent relative to the control group. Adding a deadline to either message further improves the settlement rate by another six percentage points, and increases expected revenue by 130 percent relative to the control group.

The large and positive effect of adding a deadline to the messages suggests that the weak response to the substantial price reduction in the first experiment occurs *not* because price does not matter. A plausible explanation is that, when the requested amount is lower, infringers perceive the Agency as being less serious. In other words, a lower request may not generate a higher settlement rate if it simultaneously sacrifices an implicit threat of escalation. The deadline in the messages appears to simultaneously leverage the inducement of a lower price today and signal a willingness to escalate in the future. This conclusion is further supported by heterogeneity in the effect: while the deadline has a large and positive effect on the settlement rate for relatively small firms (ten percentage points), it has no effect for larger firms (those in the top 40 percent of the firm-size distribution). This may be due to larger firms being less price-sensitive and less influenced by the implied threat, perhaps because they have more precise knowledge about the likely legal risks.

The effect of the first two messages (i.e., those without a deadline) is large and positive for all firms, and we find evidence suggesting that the acknowledgment of possible unintentionality is important for explaining this effect. At least two different mechanisms may be at work. First, the acknowledgment of unintentionality may encourage an infringer's intrinsic motivation to correct a mistake, especially given that infringement in this market is likely to be inadvertent (or indirect through third-party web designers) with small monetary claims. This effect is consistent with findings in the tax compliance literature (Feld and Frey (2002a,b)), and may be especially relevant for smaller firms. Second, the acknowledgement of unintentionality may make the settlement request conform more closely to standard invoices, providing a smoother path for processing payment, rather than initiating a lengthy internal legal review. This effect may be especially pertinent for larger firms.

Finally, the additional settlement induced by the treatment messages substantially reduces social costs. The probability that the Agency will engage a professional law firm for enforcement decreases by roughly

five to 12 percentage points (a 10-27 percent reduction) for the treatment groups relative to the control group. The estimated savings in legal costs (for the Agency and infringing firms combined) range from about three to seven percent of the Agency's in-house settlement revenue. This calculation is likely to underestimate the savings in total social costs because it does not take into account the opportunity cost of time or the stress associated with prolonged disputes.

Related literature

This paper brings together two literatures that investigate why disputes fail to reach settlement and instead go to trial. The first group of papers is from law and economics (e.g., Bebchuk (1984), Priest and Klein (1984), and the survey by Spier (2007)). Empirical work in this area studies the determinants of settlement, including the expected size of the verdict, litigation costs, asymmetric beliefs, and legal environments (e.g., Danzon and Lillard (1983), Waldfoegel (1995), and Fournier and Zuehlke (1996)). The second set of studies is from behavioral law and economics, and uses lab experiments to study the effects of psychological factors on pre-trial bargaining (see a survey by Jolls et al. (1998)). For example, Rachlinski (1996) shows that, consistent with prospect theory (Kahneman and Tversky (1979)), disputants are more likely to accept a settlement offer if they view it as a gain relative to provided references. Korobkin and Guthrie (1994) find that victims are more likely to settle if there is a sympathetic explanation for the harm they have suffered, reflecting general concerns for fairness and reciprocity (e.g., Güth et al. (1982) and Kahneman et al. (1986)). We contribute to both literatures by introducing a novel dataset examining pre-filing settlements and two large-scale field experiments that provide exogenous variation in settlement methods. Our results, collectively, show that infringers respond both to price and threat incentives that are typically captured by an economic model of settlement, as well as to motivations outside a standard cost-benefit framework.

Most of the IP enforcement literature studies patents (e.g., Lanjouw and Schankerman (2001, 2004)).⁵ Our paper contributes to this literature with empirical evidence on copyright enforcement and settlement. Recent studies in economics have examined the effects of stronger enforcement on file-sharing activities and sales in the music and movie industries (Bhattacharjee et al. (2006), Danaher et al. (2014), and Reimers (2015)). Our study differs from these papers and from the broader consumer piracy literature (see survey by Waldfoegel (2012)) in its investigation of infringement by businesses and its focus on settlement outcomes.

⁵Examples of theoretical papers on patent litigation include Meurer (1989) and Bessen and Meurer (2006), and empirical studies include the role of cross-licensing in litigation (Galasso et al. (2013)) and non-practicing entities (Bessen et al. (2011), Scott Morton and Shapiro (2014), and Cohen et al. (2014)). It is important to note that the two contexts, though conceptually similar, differ in substantial ways. First, liability in our context is generally clear, given the accuracy of detection technology and the simpler nature of copyright infringement; liability is much harder to establish in patent disputes. Second, patent claims (at least for complaints filed in court) are typically much larger. These differences may limit the effectiveness of our particular interventions in patent contexts.

In a related paper, Fellner et al. (2013) randomize messages mailed to Austrian households that evade TV licensing. They observe a significant positive effect on compliance as a result of threat, while moral appeal (emphasizing compliance as a matter of fairness) and social-norm information (highlighting the high level of compliance) have no effect. Similarly, a law literature on copyright litigation strategies argues that sanctions can be effective in inducing settlement because statutory damages for willful infringement are prohibitively high (Lantagne (2004) and DeBriyn (2012)). One salient difference between our paper and these studies is the nature of infringement. For example, Fellner et al. (2013) suggests that moral persuasion may lack effect in their setting due to a particularly dishonest sample of infringers. In contrast, we are among the first to assess infringement incidences that are likely to be inadvertent.

Finally, our paper joins the small but growing body of work that uses large-scale field experiments to study the effects of small, low-cost interventions (especially communication). Other enforcement contexts include tax compliance and enforcement (Hallsworth et al. (2014) and Perez-Truglia and Troiano (2015)) and loan repayments (Cadena and Schoar (2011) and Karlan et al. (2015)). DellaVigna and Gentzkow (2010) surveys empirical evidence on the effects of persuasive communication in several domains and defines a ‘persuasion rate’ that harmonizes results across studies.⁶ Compared to the 17 previous studies surveyed in that paper, in which the persuasion rate ranges from 0.7 to 29.7 percent, our results show generally high persuasion rates of 14 to 20 percent for messaging within the first 30 days.

The rest of the paper proceeds as follows. Section 2 provides the background and introduces the two experiments. Section 3 describes the sample, the variables, and the descriptive results. Section 4 develops a conceptual framework that potentially explains the observed effects. Section 5 reports the regression results and discusses the plausible mechanisms and the effects on social costs, and Section 6 concludes.

2 Background and Experiments

The Agency monitors the online use of images that satisfy the following criteria: (i) the image is governed by a ‘rights-managed’ license model, which allows a licensee a one-time use of the image for a fee that depends on the specified scope of use (e.g., a high-resolution image placed on a firm’s home page for three months); and (ii) the Agency distributes the image exclusively, allowing the Agency to determine whether the use is unauthorized by checking its own licensing database. Images that satisfy these criteria are the most expensive ones; they account for less than five percent of the images in the Agency’s complete portfolio. The Agency pursues enforcement only if infringement involves commercial use of an image, and it drops

⁶The notion of a ‘persuasion rate’ estimates the percentage of receivers that change behavior among those that receive a message and are not already persuaded.

cases that involve small non-profit organizations or current customers who use the images beyond the scope of their current license.

For cases it pursues, the Agency sends a letter to the infringing firm, which (1) informs the firm of the detected infringement, (2) requests the removal of the image(s) from the firm’s website, and (3) asks for a settlement to cover the past infringing use. Prior to both experiments, the requested amount was the baseline licensing fee plus an additional \$400 charge per image to recoup the Agency’s enforcement costs. The additional \$400 enforcement amount was never broken out separately in the letter. Table 1 lists the six use categories for a single image, the corresponding baseline licensing fees for the sample period we study, and the percentage of cases from our combined experimental groups (described later) that corresponds to each use category. When multiple images are involved in a case, the licensing fee is the sum for all the images. The list prices posted on the Agency’s website also vary across other dimensions, including image type (e.g., theme) and duration of use. The fee charged for an infringing use is a weighted average of a few listing prices from within one of the six categories listed in table 1.

Table 1: Licensing fee (the baseline price)

Image-use category	Licensing fee	% of cases
Low resolution & secondary page	\$380	28
High resolution & secondary page	\$475	31
Low resolution & home page	\$565	8
Low resolution & repeat pages	\$665	4
High resolution & home page	\$775	12
High resolution & repeat pages	\$825	7
Multi-image cases	Sum of fees for all images	10

Notes: This table describes the licensing-fee part of the total requested amount. The first six rows apply to single-image cases. The corresponding percentages are based on our combined experimental groups, described later.

Two experiments

2.1 The requested amount

Between October and December 2013, for cases in the U.S. and Canada, the Agency exogenously selected a subset of cases for which it removed the enforcement cost completely (the ‘Low Price’ group); for the rest of the cases, it continued to add \$400 per image on top of the licensing fee (the ‘High Price’ group). The letters to the two groups are otherwise identical. That is, no explanation is given about the price reduction for the Low Price group. We refer to this as the *pricing* experiment. This exogenous price variation is valuable because, otherwise, it is not possible to separate the effect of a price change from the way in which the image is used, as the licensing fee is determined completely by the image use category.

2.2 The messages

In January 2014, the Agency implemented two changes to its enforcement policy. First, the Agency no longer requested any settlement amount from “less valuable” infringers, although it still requested removal of the images from the infringer’s website. More precisely, infringers are categorized into five ‘tiers’ using a formula from Dun & Bradstreet, with tier 1 being the most likely to settle and tier 5 the least likely; the Agency no longer pursues settlement from firms in tiers 4 and 5.⁷ Second, the Agency no longer added any enforcement cost to any cases. Following these policy changes, we collaborated with the Agency to conduct a second experiment, also for infringement cases in the U.S. and Canada.

In the second, *message* experiment, we provide information about the price reduction while maintaining the low price (i.e., with no additional \$400 charge). Cases are allocated into four treatment groups and one control group. All groups are requested to pay the same licensing fee (given the image use), although the letter’s wording varies. The wording used for the control group is the same as in the pricing experiment; it provides no information about the price reduction. For the treatment groups, one of four extra messages (in boldface) is placed just below the invoiced amount. Figure 1 shows the positioning of the message in the letter. The four messages differ by how the forgiven amount is presented, and whether a 22-day deadline is imposed. As justification for the price reduction, all four messages acknowledge that the infringement may have been unintentional. The four messages are:


1. (Control group) no extra message.
2. (Waive the Cost) Please note, we are only charging the average licensing fee for commercial use of the rights-managed image(s) found on your website. The Agency has incurred additional costs of \$400 per image related to the pursuit of this matter; we are currently waiving this cost, as we understand this unlicensed use may have been unintentional.
3. (Waive the Cost + Deadline) Please note, we are only charging the average licensing fee for commercial use of the rights-managed image(s) found on your website. The Agency has incurred additional costs of \$400 per image related to the pursuit of this matter; we are currently waiving this cost, as we understand this unlicensed use may have been unintentional. However, cases not settled before xx-xx-xx will be charged the additional \$400 cost per image related to the pursuit of this matter.
4. (Discount) Please note, this settlement amount is \$400 per image lower than the average unauthorized use settlement offer. We are providing this discount, as we understand this unlicensed use may have been unintentional.

⁷We do not observe the formula that maps firm observables to tiers, but firms with lower annual revenues tend to have a higher tier number.

5. (Discount + Deadline) Please note, this settlement amount is \$400 per image lower than the average unauthorized use settlement offer. We are providing this discount, as we understand this unlicensed use may have been unintentional. However, cases not settled before xx-xx-xx will no longer be eligible for this discount.

Figure 1: Settlement demand letter example

Details

	<p>██████████ Beech tree through four seasons (digital Composite)</p> <p>Collection: ██████████</p> <p>Photographer: ██████████</p> <p>Usage: Web - Corporate or promotional site</p> <p>Placement: Repeated icon</p> <p>Territory: United States</p> <p>Release Information: No release, but release may not be required.</p>	<p>Image price: \$825.00</p>
---	--	---

Notes: This settlement demand does not represent an approval on the part of ██████████ for the unauthorized use or uses of this image identified to date and referenced herein. The payment of this settlement demand together with your immediate cessation of use of the image (unless you have separately licensed this image for future use), including uses by your owners, directors, employees, agents, clients and/or licensees, in any and all media, will release you from any legal claims by ██████████ relating solely to this identified past infringement.

Total

Subtotal:	\$825.00 USD
Tax	\$0.00 USD
Total Due Upon Receipt	\$825.00 USD

Please note, this settlement amount is \$400 per image lower than the average unauthorized use offer. We are providing this discount, as we understand this unlicensed use may have been unintentional.

Notes: A letter comprises seven pages in total, with the invoiced amount displayed on the third page. All letters in the experiments are of the same format except for the extra message for the four treatment groups in the message experiment (as illustrated in this example).

For both experiments, an infringer could contact the Agency to try to negotiate a lower price. Employees in the Agency’s compliance department, “case handlers,” respond to calls from infringers and have latitude to negotiate with them as they see fit. According to the Agency, case handlers are not instructed to treat groups differently in either experimental period, mainly so as not to confuse the handlers. In principle, the handlers can observe the group to which a particular case is allocated. Our interviews with the handlers, however, suggest that they tend not to incorporate that information during conversations with infringers.

2.3 Allocation process

The Agency allocated cases to different groups weekly during each experiment. For a given week, some cases were assigned as “trial cases,” and all other cases were assigned to a “default” group. The “default” group is the High Price group in the pricing experiment, and it is the control group in the message experiment. After cases were allocated to the trial and default groups, cases in the trial group from each combina-

tion of infringer tier and price category were placed sequentially into each treatment arm.⁸ The number of cases allocated as trial cases each week depended on the Agency's overall workload. In the pricing experiment, we use data from five weeks during which trials were conducted. In the message experiment, trials were conducted every week between the end of January and the end of March 2014.

It is not ideal that the allocation procedure did not follow a strict randomization protocol. In particular, we do not know exactly how cases were assigned to the trial and default groups in the first step. According to the Agency, trial cases were selected sequentially from a list of infringements generated each week by the detection software, which had been screened to eliminate authorized uses. The Agency is not aware of any dependence between the ordering of cases generated by this process and settlement likelihood. This is reasonable because neither the detection technology nor the Agency's internal authorized-use screening team has information on the potential settlement outcomes. Nevertheless, to mitigate this issue, we conduct balance tests on the requested amount and infringer characteristics in order to verify that the treatment and control groups are consistent with random allocation (see table 3 below).

3 Sample, variables and descriptive results

Our dataset includes all the infringement instances for which the Agency pursued settlement during the two experiments. The information includes the initial requested amount, characteristics of infringing use, and settlement outcomes. The data on infringer characteristics come from Dun & Bradstreet and include annual sales, number of employees, age, industry, and firm location.

The final sample excludes cases as defined here. First, because the Agency stops pursuing settlement from infringers in tiers 4 and 5 during the message experiment, we also remove these cases in the pricing experiment, in order to make the populations of the two experiments comparable. This excludes 30 percent of all cases, leaving 5,660 cases. We further exclude cases for which (1) the mail is undeliverable (ten percent of the 5,660 cases); (2) the case is later closed because the infringer is able to show a valid license or permission from the photographer (6.3 percent); (3) there is no information on the firm's annual sales (8.5 percent); or (4) more than six images are involved, ensuring that the maximum number of images in all groups is the same (19 observations).⁹ The final sample includes a total of 4,278 cases, of which 1,983 come from the pricing experiment and 2,295 from the message experiment.

Panel 1 in table 2 summarizes the requested amount, inclusive of the additional \$400 charge when

⁸In the pricing experiment, there was also a small treatment arm in which a \$100 enforcement cost was added to the licensing fee. The Agency discontinued this price point after a short time. We exclude these observations in the current results, but including them does not qualitatively change the basic conclusions.

⁹Our basic conclusions do not change because of any of the exclusions. Table A6 in the Online Appendix reports results using unfiltered data (that is, without dropping cases for which the mail is returned or sales information is missing, etc).

applicable, and the number of images involved in each case. The median requested amount is \$699. Ninety percent of all cases involve a single image, and nearly 60 percent are in the two lowest-priced use categories (as shown in table 1). For cases involving multiple images, the median number of images is two.

Table 2: Summary statistics of case and firm characteristics

Variables	N	Mean	SD	Min	Median	Max
<u>Panel 1: Price and image use</u>						
Requested amount	4,278	748.0	438.3	354.9	699.0	5040.0
Single-image case	4,278	0.90	0.30	0	1	1
Image count (if a multi-image case)	437	2.49	0.85	2	2	6
<u>Panel 2: Firm characteristics</u>						
Annual sales (\$000's)	4,278	23165.7	423542.4	0.3	320.0	23.7M
Employees	4,278	112.9	1348.5	0	5	67800
Age	4,155	17.8	21.1	0	11	271

Notes: This table summarizes the requested amount, the number of images, and the characteristics of the firms, using cases in both experiments. Firm characteristics are provided by Dun & Bradstreet based on identifying information of the infringing firm (such as name and address). Industry and state (province) dummies are omitted.

Panel 2 in table 2 summarizes infringer size and age. The median firm has \$320K in annual sales; the median number of employees is five, and the median firm age is 11 years. The distribution of several firm characteristics (e.g., annual sales) is highly skewed. Annual sales and the number of employees are strongly correlated (the correlation is 0.92), so we use only the former to measure firm size. The top ten industries (defined by two-digit SIC codes and not shown in table 2) are responsible for 63 percent of all cases, and they are largely service industries. The three states with the largest number of infringers are California, Florida, and New York.

Table 3 compares key pre-treatment variables for different groups for the two experiments. Because sales and age are skewed, we summarize both the mean and the median of the variables and the mean of their log transformation. Panel (a) shows that, in the pricing experiment, there is no significant difference in the average baseline price (i.e., the licensing-fee portion of the requested amount) or the summary statistics of the distributions of firm size and age between the High Price and Low Price groups. Panel (b) summarizes the same information for the control and four treatment groups of the message experiment, along with p-values of two-sided t-tests (or median tests) comparing each treatment group to the control group. The variables are similar between the treatment groups and the control, with the following exceptions: compared to the control, (1) the average requested amount in the ‘Waive the Cost’ group is higher; (2) firms in the two deadline message groups are smaller in log(Sales), even though the mean and the median of the level of sales are not statistically different from the control; and (3) firms in the ‘Waive the Cost + Deadline’ group

Table 3: Balance check

(a) Pricing experiment

Group	N	Baseline price	Sales, \$1,000		log(Sales, \$)	Age		log(Age+1)
			mean	median	mean	mean	median	mean
Low Price	707	642.6	15533	324	13.1	18	11	2.5
High Price	1,276	623.2	13797	281	13.0	17	11	2.5
(p-value)		(0.22)	(0.84)	(0.44)	(0.41)	(0.42)	(0.58)	(0.92)

(b) Message experiment

Group (p-value, diff. to control)	N	Baseline price (=Requested amount)	Sales, \$1,000		log(Sales, \$)	Age		log(Age+1)
			mean	median	mean	mean	median	mean
Control	598	590.7	30614	390	13.4	19	11	2.6
Waive the Cost	378	636.8 (0.03)	50682 (0.40)	385 (0.91)	13.4 (0.93)	19 (0.86)	11 (0.83)	2.5 (0.86)
Waive the Cost + Deadline	414	595.3 (0.82)	7951 (0.16)	290 (0.13)	13.1 (0.03)	16 (0.04)	9 (0.17)	2.5 (0.06)
Discount	448	588.4 (0.90)	10112 (0.19)	305 (0.13)	13.2 (0.13)	18 (0.48)	12 (0.73)	2.5 (0.85)
Discount + Deadline	457	580.2 (0.59)	55206 (0.61)	360 (0.47)	13.0 (0.01)	17 (0.18)	11 (0.59)	2.5 (0.37)

Notes: Panel (a) summarizes the baseline price (i.e., the licensing-fee portion of the total requested amount), the number of images, and firms' basic characteristics between the two groups in the pricing experiment. Two-sided t-test p-values are reported in parentheses. Panel (b) summarizes the same information by group for the message experiment. Two-sided t-test (or median test) p-values are between each treatment group and the control group.

are younger.

Consultation with the compliance team at the Agency does not identify any sources of these incidences of imbalance. If the suggestions of imbalance here are not false positives, they are likely to make our estimates of the treatment effects more conservative. This is because our results consistently show that settlement is less likely for smaller and younger firms and for cases with a higher requested amount. We further verify that the regression results without using controls (reported in table A1 in the Online Appendix) are similar to those using controls. Both of these are also consistent with the raw data, indicating that the allocation is not systematically correlated with observable variables.

About 30 days after sending the first letter, the Agency sends follow-up letters unless the case has already been settled. Eventually, cases that do not settle are sent to external law firms. For cases that settle within our sample period (all cases are older than five months), 51 percent settle by 30 days, and 98 percent settle by 120 days. 90 percent of the cases that are sent to external law firms are sent between 58 and 159 days after the first letter. In the main analysis, we focus on the 30-day outcomes as a clean assessment of the interventions. Two factors confound assessments made in later time periods. First, more follow-up letters are sent to groups that have lower initial settlement rates. These endogenously-determined follow-up

actions are designed to encourage settlement and should cause the worst-performing groups to ‘catch up’ over the longer horizon. Second, software constraints at the Agency accidentally resulted in substantially more intense follow-up for groups designated as the ‘default’ group in each experiment.¹⁰ In section 5, we provide estimates that bound the longer-term outcomes, such as the effects of the interventions on the probability of escalation to an external law firm.

Table 4 summarizes 30-day outcomes. Taking the settlement probability into account, the average revenue per case is \$104. For settled cases, the average revenue collected is \$555. Infringers negotiate a lower settlement amount for 38 percent of settled cases, and for those cases, the median reduction is 36 percent. The overall settlement rate is 19 percent. Broken down by firms’ responses to the initial letter, seven percent settle directly without contacting the Agency; 27 percent contact the Agency (43 percent of which settle by 30 days); and 66 percent ignore the letter. Our analysis pools the two types of settlements together (that is, direct settlement without contacting the Agency and settlement after making contact).

Table 4: 30-day outcomes

(a) 30-day revenues

Variables	N	Mean	SD	Min	Median	Max
Revenue per case	4,278	104.2	255.8	0.0	0.0	2992.3
Revenue (if settled)	803	555.2	313.4	91.0	475.0	2992.3
Paid a reduced amount? (if settled)	803	0.38	0.49	0	0	1
Negotiated reduction % (if reduced)	305	0.36	0.17	0.02	0.35	0.88

(b) 30-day settlement rate and infringer response

Contact/Outcome	Settle (19%)	Not Settle (81%)
Contact Agency (27%)	12%	15%
Do not contact Agency (73%)	7%	66%

Notes: Panel (a) summarizes the 30-day revenues using cases in both experiments. Revenue is coded as \$0 if the case does not settle within 30 days. Panel (b) summarizes the 30-day settlement rate and detailed responses by infringers.

3.1 Descriptive results

Table 5 compares the settlement outcomes for the two price groups in the pricing experiment. On average, the 30-day settlement rate of the Low Price group is three percentage points higher than that of the High Price group (15.6 vs. 12.5 percent with a p-value of 0.054). Although this represents a 25 percent increase, the difference is only marginally significant. For settled cases, the average revenue collected is \$211 lower in the Low Price group. This suggests that more extensively negotiated reductions for the higher-priced group

¹⁰Recall from Section 2.3 that the ‘default’ group was the High Price group in the pricing experiment and the control group in the message experiment. We provide a detailed description of the Agency’s follow-up actions in Section C of the Online Appendix.

offset half of the difference in the requested amount. Despite the higher settlement rate, the expected 30-day revenue of a case is \$10 lower for the Low Price group, but the difference is not statistically significant.

Table 5: Price effects on settlement outcomes (raw data)

	Low Price	High Price	Difference	p-value
<i>All cases:</i>				
Settlement probability	0.156	0.125	-0.031	0.054
Revenue	82.6	92.4	9.8	0.432
N	707	1,276		
<i>Settled cases:</i>				
Revenue	530.9	741.5	210.6	0.000
Whether discounted	0.36	0.64	0.28	0.000
Discount	81.9	284.4	202.5	0.000
N	110	159		

Notes: This table summarizes the 30-day outcomes of the High Price and the Low Price groups in the pricing experiment.

Table 6 reports the settlement probability for the control and treatment groups in the message experiment. Table 6a shows that, for an average case, relative to the control group, adding either of the two non-deadline messages increases the settlement probability by about nine percentage points (a 64 percent increase); adding either of the two deadline messages improves the 30-day settlement probability by about 15 percentage points (a 107 percent increase relative to the control group); and all the improvements are highly statistically significant.¹¹ For each of the two different ways of presenting the forgiven amount (i.e., “Waive-the-Cost” and “Discount”), the difference between the deadline and non-deadline versions is about five or six percentage points (p-values are 0.11 and 0.05, table 6b). However, given that a deadline is added (or not), there is no significant difference between the two different presentations of the forgiven amount. As shown in later regression results (table 9), the messages’ effects on expected revenue remain consistent with those for the settlement probability. This is because the requested amount is the same across different groups given the same use category, and the revenue collected per settled case is similar across groups.

The two experiments combined show that, on average, a substantial reduction in the requested amount generates a small increase in the settlement rate (and, consequently, a small and statistically insignificant difference in expected revenue). In contrast, for the same reduced request, including any of the treatment messages generates a large and statistically significant increase in the settlement rate; and, holding the rest of the message content constant, adding a deadline alone further increases the response by a significant amount. Although the estimates are obtained from different samples, the two sample periods are close in

¹¹We report the 30-day outcomes for the deadline messages for ease of comparison. Pre-deadline (22-day) outcomes are reported in a later section and, in summary, show that the two deadline messages improve both the settlement rate and the expected revenue by about 150 percent relative to the control group (0.25 versus 0.10, and \$120 versus \$47).

Table 6: Message effects on settlement probability (raw data)

(a) Comparing the treatment groups to the control

Group	N	Settlement probability	Difference to Control	(p-value) diff. to control
Control	598	0.14		
Waive the Cost	378	0.24	0.09	(0.00)
Waive the Cost + Deadline	414	0.29	0.14	(0.00)
Discount	448	0.24	0.10	(0.00)
Discount + Deadline	457	0.30	0.16	(0.00)

(b) Comparing between different treatment groups

	No deadline	With deadline	Effect of a deadline	(p-value)
Waive the Cost	0.24	0.29	0.05	(0.113)
Discount	0.24	0.3	0.06	(0.046)
Difference in presenting the forgiven amount (p-value)	-0.003 (0.91)	-0.01 (0.68)		

Notes: (A) Summary of 30-day settlement probabilities for the four treatment groups and the control group in the message experiment. Two-sided t-test p-values are between each treatment group and the control group. (B) Comparison of the settlement probabilities between different treatment groups and p-values for these differences.

time, and the settlement rate of the Low Price group in the pricing experiment and that of the control group in the message experiment (both add no additional charge and include no extra message) are not statistically different from each other (0.156 vs. 0.141 with a p-value of 0.44).

4 Conceptual Framework

In this section, we propose two mechanisms that, taken together, help to explain the findings described in Section 3.1. First, we develop a simple rational model in which the infringer trades off the disutility of paying to settle and the expected liability risk of resisting settlement. The infringer begins with limited information about what may happen if he does not settle immediately. These beliefs are exogenously influenced by the experiments' price reduction or messaging (especially through the provision of information about the price of settling). We show that this 'expected-payoff' mechanism alone is sufficient to capture the basic results of both experiments.

Apart from providing relevant price information, the treatment messages also contain an acknowledgment of possible unintentionality. We provide evidence suggesting that unintentional infringement is prevalent in our context. Given this nature of infringement and the fact that the monetary stakes are not large, this acknowledgment may encourage more collaborative outcomes by shifting the frame of decision-making for some infringers away from a pure cost-benefit analysis (as captured in the 'expected-payoff' mechanism).

We refer to this mechanism as a ‘collaborative-mindset’ mechanism. Within the message experiment, because the two non-deadline messages provide price information and acknowledge possible unintentionality, while the control does neither, the difference between them may be the result of both mechanisms. In the last part of this section, we discuss possible ways to use heterogeneity in the data to clarify whether or not the ‘collaborative-mindset’ mechanism plays an important role in encouraging settlement.

4.1 Expected-payoff mechanism

We present a two-period, single-agent decision model with incomplete information, in which the infringer decides whether or not to settle in each period. The settlement offers for the two periods are, respectively, p_1 and p_2 ; the amount requested for delayed settlement is always at least as high as that for immediate settlement (i.e., $p_2 \geq p_1$). In period 1, the infringer chooses whether to settle or to wait until period 2. If he settles, he pays the requested amount p_1 . If he waits until period 2, he calculates an expected second-period payoff as a function of (i) his expectation of the probability of enforcement in period 2, and (ii) his expected payoff, conditional on enforcement.¹² We solve for the second component first.

Conditional on enforcement, an infringer’s second-period payoff is

$$-\pi_2^i = \begin{cases} -p_2 & \text{if settle in period 2} \\ -\varepsilon_2 & \text{otherwise,} \end{cases}$$

where $\varepsilon_2 > 0$ is the expected cost to the infringer if he does not settle in period 2. This cost could include, for example, legal expenses, opportunity costs of time and stress, and reputational costs of being engaged in a prolonged dispute. The infringer observes ε_2 in period 2 and, thus, settles when $-p_2 > -\varepsilon_2$. We denote the infringer’s second-period payoff when the Agency enforces as $-\pi_2^i(p_2) = -p_2 I(\varepsilon_2 > p_2) - \varepsilon_2 I(\varepsilon_2 \leq p_2)$.

In period 1, the infringer knows the distribution of ε_2 , which is denoted as $H(\cdot)$. Thus, the infringer’s expected second-period payoff, calculated in period 1 and conditional on enforcement, can be written as

$$-E[\pi_2^i(p_2)] = -(1 - H(p_2))p_2 - H(p_2)E_{\varepsilon_2}[\varepsilon_2 | \varepsilon_2 \leq p_2].$$

Note that $E[\pi_2^i(p_2)] > 0$ (i.e., it is costly to the infringer when the Agency chooses to enforce in period 2), and that $E[\pi_2^i(p_2)]$ is increasing in p_2 (i.e., it is more costly to the infringer when the requested settlement amount is higher), all else equal.¹³

¹²We assume that the infringer ignores the possibility that the Agency may be willing to negotiate to a different settlement offer at a later stage.

¹³ $E[\pi_2^i(p_2)]' = \frac{\partial [p_2 - H(p_2)p_2 + \int_0^{p_2} \varepsilon_2 h(\varepsilon_2) d\varepsilon_2]}{\partial p_2} = 1 - H(p_2) \in (0, 1)$.

Next, we turn to the infringer's expectation of the likelihood of enforcement in period 2. Motivated by the context of our empirical application, we assume that an infringer's expectation of enforcement depends on his expectation of the Agency's payoff from enforcing. Without precise knowledge about this payoff, the infringer may use available price information, and in particular, the price for delayed payment, p_2 , to infer the likelihood of Agency enforcement. Denote the infringer's expectation of the Agency's payoff from enforcing as $v(p_2) - \varepsilon_a$, where $\varepsilon_a \sim G(\cdot)$ is the Agency's cost of pursuing settlement. We assume that $v'(p_2) > 0$ (i.e., the Agency's payoff from enforcement increases with p_2). The infringer's expectation of the Agency's second-period payoff is:

$$\pi_2^a = \begin{cases} v(p_2) - \varepsilon_a & \text{if enforce} \\ 0 & \text{otherwise.} \end{cases}$$

Thus, the infringer's expectation of the likelihood of enforcement is $G(v(p_2))$. This implies that the probability of enforcement increases in p_2 .¹⁴

Thus, in period 1, the infringer's payoff function can be written as:

$$-\pi_1^i = \begin{cases} -p_1 & \text{if settle in period 1} \\ -\beta G(v(p_2))E[\pi_2^i(p_2)] + \varepsilon_1 & \text{otherwise,} \end{cases} \quad (1)$$

where β is a standard discount factor; $G(v(p_2))$ is the infringer's expectation of the probability of enforcement in the second period; $E[\pi_2^i(p_2)]$ is his expected second-period payoff conditional on enforcement; and $\varepsilon_1 \sim F(\cdot)$ reflects an infringer's private cost of waiting until period 2. This defines the share of infringers that settle in period 1, s_1 , as

$$s_1 = F(-p_1 + \beta G(v(p_2))E[\pi_2^i(p_2)]). \quad (2)$$

We are interested in the implications of the two experiments for the share of infringers that settle in period 1. The experiments vary either the actual price (the pricing experiment) or the information disclosed to an infringer about price (the message experiment). In the pricing experiment, all infringers face a single settlement offer (either the High Price or the Low Price). Because this offer is the only available price information, the infringer may associate it with the value of pursuing the case from the Agency's point of view and with the amount he may need to pay if he delays settlement until period 2. In other words, $p_2 = p_1 = p$ for infringers in both groups. To obtain the expected treatment effect of a higher settlement

¹⁴Equivalently, $\partial G(v(p_2))/\partial p_2 = g(v(p_2))v'(p_2) > 0$.

offer, replace p_2 and p_1 with p in equation (2), and take the derivative with respect to p :

$$\frac{\partial s_1}{\partial p} = f(-p + \beta G(v(p))E[\pi_2^i(p)]) \left(\underbrace{-1 + \beta G(v(p))E[\pi_2^i(p)]'}_{\text{traditional price effect (-)}} + \underbrace{\beta g(v(p))v'(p)E[\pi_2^i(p)]}_{\text{threat effect (+)}} \right). \quad (3)$$

Equation (3) illustrates the fact that a higher settlement offer may simultaneously encourage and discourage immediate settlement. On the one hand, the ‘traditional price effect’ implies that an increase in the offered price should decrease the probability of immediate settlement because the demand for settling is downward sloping (i.e., $-1 + \beta G(v(p))E[\pi_2^i(p)]' < 0$).¹⁵ On the other hand, a higher price may suggest a greater ‘threat’ of escalation (i.e., $\beta g(v(p))v'(p)E[\pi_2^i(p)] > 0$), which, in turn, encourages immediate settlement. These two effects offset each other and, thus, may explain why a substantial reduction in price generates only a small difference in the observed settlement rate.

In the message experiment, the control group continues to face a single price. In contrast, infringers in the treatment groups receive information about two prices: they receive the same settlement offer as the control group and, at the same time, they are informed that the offer is \$400 lower than it could have been. In other words, all four treatment groups face the same p_1 as the control, but they receive a higher p_2 . The second, higher price may lead infringers to infer a higher recoupable revenue by the Agency and, hence, a greater threat of further enforcement and a higher amount due for delayed settlement.¹⁶ To obtain the expected effect of the treatment messages relative to the control group, consider the derivative of equation (2) with respect to p_2 :

$$\frac{\partial s_1}{\partial p_2} = f(-p_1 + \beta G(v(p_2))E[\pi_2^i(p_2)]) \left(\underbrace{\beta G(v(p_2))E[\pi_2^i(p_2)]'}_{\text{intertemporal price effect (+)}} + \underbrace{\beta g(v(p_2))v'(p_2)E[\pi_2^i(p_2)]}_{\text{threat effect (+)}} \right). \quad (4)$$

Both terms in equation (4) are positive. In contrast to the pricing experiment, the ‘intertemporal price effect’ and the ‘threat effect’ of a higher p_2 are synergistic (rather than oppositional); both effects make ‘settling in period 1’ a more attractive option. This can explain why the messages generate a large and positive response from infringers. Furthermore, relative to the messages without a deadline, the two messages with a deadline should yield a stronger result, because both the intertemporal price trade-off and the increased threat are more explicit.

¹⁵Even though a higher settlement offer also means a higher price in period 2, the direct effect on the disutility of paying a greater price in period 1 dominates; that is, $-1 + \beta G(v(p))E[\pi_2^i(p)]' = -1 + \beta G(v(p))(1 - H(p)) < 0$.

¹⁶Even infringers who receive the treatment messages that do not include an explicit deadline may infer that the discounted offer will no longer apply to a delayed settlement. This is consistent with findings in the marketing literature, in which researchers have found that placing a sale sign on an item is sufficient to increase demand for the item even without changing the actual price, because consumers expect the price to go up after the promotion period (e.g., Inman et al. (1990) and Anderson and Simester (1998)).

To summarize, pricing not only affects the infringer’s ‘demand for settlement’ directly, but also influences his belief about his potential future liability risk. According to the model, the reason that a substantially lower settlement offer alone has little effect on the settlement rate is *not* because price does not matter. Rather, it is because infringers perceive the Agency as being less serious about enforcement when the requested amount is lower. In other words, a lower request may not generate a higher settlement rate if it simultaneously sacrifices an implicit threat of escalation. Decoupling the two prices, as the treatment messages do in the second experiment, allows the Agency to simultaneously leverage the inducement of a lower price today and the threat of future escalation.

4.2 Collaborative-mindset mechanism

The expected-payoff mechanism modeled in the previous section explains the small price effect from the pricing experiment, which is cleanly identified in the data. The incremental effect of adding a deadline to the treatment message is also cleanly identified because comparisons between the deadline and non-deadline messages control for the rest of the message content. The expected-payoff mechanism is also sufficient to explain this effect, as an explicit deadline strengthens both the intertemporal price trade-off and the threat. However, within the message experiment, the non-deadline messages contain two pieces of information that the control lacks: (i) pricing information, and (ii) acknowledgment of possible unintentionality. Thus, the large and positive effect of these non-deadline messages may be the result of both factors.

Table 7 shows that accidental or indirect infringement is prevalent in our data. The table summarizes the (self-reported) reasons for infringement, which are collected from correspondence records between the Agency and infringers that made contact during the message experiment. In 43 percent of the cases, the infringer claims to be unaware of any infringement, or to have obtained the image online, thinking that it was free to use. In 42 percent of the cases, an outside party is involved in designing the infringer’s website.¹⁷ In four percent of the cases, the firm uses a previously-licensed image outside the specified scope. While we acknowledge that the incentive for an infringer to claim unawareness may be strong, the extent of unintentional infringement (from the infringing firm’s point of view) still seems substantial. Table 7 also shows that this phenomenon is similar across infringers of all sizes.

Given the nature of these incidences, infringers may feel confused and frustrated when faced with enforcement, because they were not aware of the mistake. It is, thus, plausible that acknowledging the possible unintentionality encourages settlement by shifting the infringer’s frame of mind away from calculating le-

¹⁷We define third-party involvement to include cases for which either: (1) the party who contacts the Agency is a third-party designer, family or friend (17%); or (2) the end-user claims that a third party designed the website (25%). It is possible that some family or friends make contact because they have a legal background. However, we group these calls with third-party designers because anecdotal evidence suggests that friends and family who contact the Agency are usually involved in the website design.

Table 7: (Self-reported) reason(s) for infringement

Infringer-size quintile	N	Third-party involved	Unaware/obtain from the Internet	Improper use of license	Other reasons	Deny infringement
1	154	0.41	0.47	0.01	0.06	0.05
2	168	0.42	0.43	0.01	0.05	0.08
3	181	0.43	0.44	0.02	0.07	0.04
4	143	0.46	0.40	0.03	0.06	0.05
5	176	0.37	0.45	0.09	0.07	0.02
Total	909	0.42	0.43	0.04	0.07	0.05
Corr. with log(sales)		-0.04	0.02	0.17	-0.02	-0.06
p-value		(0.24)	(0.67)	0.00	(0.65)	(0.07)

Notes: The data are taken from correspondence records between the Agency and firms in the message-experiment sample that contact the Agency. We include infringers that make contact at *any* time during the sample period, not just in the first 30 days. There is no significant difference in annual sales, age or the number of employees between infringers that make contact (about half of the cases) and those that either ignore the letters or settle without making contact. Twenty percent of the cases for which firms make contact do not contain enough information and, thus, are not used in the tabulation. Size quintiles are based on annual sales, with Q1 being the smallest firms and Q5 the largest.

gal risks, which are likely to be small in such circumstances. At least two different sub-mechanisms may be at work. First, this acknowledgement may uphold an infringer’s intrinsic motivation to correct a mistake, which is especially effective when the amount of the claim is small. The literature on tax compliance has similar findings. For example, Feld and Frey (2002a,b) argue that paying tax is quasi-voluntary, and they show that respectable treatment by the tax authorities (i.e., a presumption of innocence) dampens tax evasion.¹⁸ The effectiveness of this acknowledgement is also consistent with findings in the bargaining literature that many people are motivated by concerns for fairness and reciprocity (e.g., Kahneman et al. (1986), Fehr et al. (1993), and Roth (1995)). Second, the acknowledgement of unintentionality may simply make the settlement request letter conform more closely to other standard invoices. For firms that have routines in place for paying bills or resolving disputes, this may provide a smoother path for processing payment, rather than invoking a costly internal legal review.

4.3 Distinguishing the two mechanisms

In this section, we discuss possible ways to utilize heterogeneity in the data to help distinguish between the expected-payoff and collaborative-mindset mechanisms in explaining the difference between the non-deadline messages and the control in the message experiment. In particular, we focus on two sources of variation that generate differential responses to these messages across infringers when (i) information about pricing or (ii) acknowledgement of unintentionality is likely to be the major driver. As we cannot directly

¹⁸Frey (1997) provides theoretical basis of psychological tax contract based on the interactions between incentives and the intrinsic motivation to pay taxes (often called ‘tax morale’). Alm and Torgler (2011) and Luttmer and Singhal (2014) survey empirical evidence for tax morale.

measure infringers' price sensitivity or their feelings about the fairness of the settlement request, the results in this section use proxies for these underlying primitives. Our ability to convincingly distinguish between these mechanisms is ultimately limited by the strength of these proxies. This, of course, is not a concern for establishing the importance of the expected-payoff mechanism in explaining the infringer's behavior as a whole since that result derives from a clean comparison of (i) the average outcomes of the two experiments, and (ii) the incremental effect of adding a deadline to the treatment messages.

In our first comparison, we consider the size of the forgiven amount, which is always \$400 per image, as a percentage of the baseline licensing fee. Baseline licensing fees range from \$380 to \$825 per image, so the relative size of the forgiven amount, $(p_2 - p_1)/p_1$, varies substantially, ranging from 43 to 107 percent. One may expect that both the intertemporal price effect and the threat effect in equation (4) are greater when the percentage difference between p_2 and p_1 is larger. If so, one may expect the effect of the non-deadline messages to increase with the percentage difference in price when the pricing information contained in the messages is the primary driver of the result (through the expected-payoff mechanism).

In our second comparison, we examine how the effect of the non-deadline messages varies across firms of different sizes. As explained previously, firm size varies substantially in our data. Consistent with many studies of demand estimation, one might expect the 'intertemporal price effect' to be smaller for larger firms, as these firms have lower borrowing costs and are typically thought to be less price-sensitive. One may also expect the 'threat effect' to be smaller for larger firms, as these firms are likely to have more information than smaller firms about underlying litigation risks. For example, additional legal expertise may lead larger firms to be less affected by small changes in communication.¹⁹ Furthermore, larger firms may have greater resources to fight a legal battle if necessary and, hence, may be less sensitive to an increased threat. Therefore, if the expected-payoff mechanism is the major driver of infringers' responses to the non-deadline messages, one may expect the size of the effect to decrease with firm size.

In contrast, if the effect of the two non-deadline messages is driven mainly by the acknowledgment of possible unintentionality, there is, a priori, no obvious relationship between the size of the effect and the relative size of the forgiven amount. Similarly, it is also not obvious how the effect should vary by firm size. Different-sized firms may react to this acknowledgment for different reasons. For example, smaller firms may be more likely to respond to an intrinsic motivation to pay, while larger firms may be more likely to settle due to the fact that this acknowledgement of unintentionality provides a smoother internal procedure for processing payment. Furthermore, regardless of the reason, the magnitude of the response may be subject to the infringer's financial capability. Thus, even if larger firms are less susceptible to small communication

¹⁹Correspondence data from the Agency show that the larger is the infringing firm, the more likely it is that the person who contacts the Agency is an internal lawyer.

changes, they may be more likely to respond because they face fewer binding budget constraints.

Finally, we expect the same heterogeneity discussed above (when information about pricing is the main driver) to also apply to the pure effect of including a deadline, because the underlying mechanism is similar.²⁰ Therefore, we can use the effect of a deadline as a benchmark, because it can be cleanly identified in the data. In other words, if the price information is the main driver of the effect of the two non-deadline messages relative to the control, we expect their subsample results to be consistent with those from comparing the messages with and without a deadline.

5 Regression results

In this section, we first present the regression results, controlling for firm and case characteristics, for an average case for both experiments. These results are consistent with the descriptive results in Section 3.1. We then report the heterogeneity (or lack thereof) in the effects across two dimensions: (1) the size of the infringing firm and (2) the extent of the price variation (in percentage terms) that is either induced by the pricing experiment, or implied by the message experiment. We follow the results with a discussion of the underlying mechanisms in light of the conceptual framework developed in Section 4. Finally, we report the effects of the treatment messages on legal costs.

5.1 Results for an average case

Pricing experiment

Table 8 provides OLS and logit regression analyses of the effect of charging the High Price (relative to the Low Price group) on the likelihood of settlement and revenue. The control variables in the regressions include the baseline licensing fee; characteristics of the infringing use, including image-use category indicators and number of images; characteristics of the infringer, including $\log(\text{annual sales})$, $\log(\text{age} + 1)$, ten ‘top’ industry indicators (defined by two-digit SIC codes), a U.S. dummy, and ten ‘top’ state/province indicators; and week dummies indicating when the settlement letter was mailed.²¹ For all regressions throughout the paper, we use the same set of control variables and report robust standard errors, clustered by industry.

The marginal effects of the High Price dummy are consistent with the raw data (table 5). On average, firms appear to be not price-elastic, in the sense that the substantial reduction in the requested amount does not result in higher expected revenue. Firms are more likely to settle if they are offered a lower price (the

²⁰That is, adding a deadline simply strengthens the inter-temporal price and the threat effects relative to the two messages without a deadline.

²¹Seventy-six firms (four percent) do not have age information. To preserve the sample size, a dummy variable indicating that the age information is missing is included in the regressions.

Table 8: Price effects on settlement outcomes

Dependent variable	Settle	Revenue for settled cases	
	Logit (1)	OLS (2)	OLS (3)
High Price dummy	-0.03* (0.02)	255.84*** (34.79)	13.72 (11.45)
Use cat 2: High res. + secondary page	0.02 (0.02)	-131.03 (80.16)	3.65 (27.71)
Use cat 3: Low res. + home page	-0.03 (0.05)	-325.96** (143.73)	-24.43 (49.89)
Use cat 4: Low res. + repeat pages	0.04 (0.06)	-285.46 (178.13)	30.10 (74.11)
Use cat 5: High res. + home page	-0.00 (0.07)	-371.27 (249.95)	-12.47 (94.20)
Use cat 6: High res. + repeat pages	0.00 (0.07)	-366.29 (263.55)	-4.53 (93.50)
Use cat 7: Multi-image cases	0.03 (0.08)	-557.20*** (190.23)	20.58 (99.43)
log(Baseline price)	-0.04 (0.09)	893.66*** (319.06)	41.77 (124.66)
Number of images	-0.01 (0.04)	346.52** (156.38)	17.09 (75.92)
log(Annual sales)	0.01*** (0.00)	23.07* (12.63)	17.20*** (3.38)
log(Age+1)	0.02** (0.01)	-7.09 (13.25)	10.03 (6.75)
Missing-age dummy	0.10** (0.04)	-97.38 (99.54)	46.87 (36.44)
Top 10 industry dummies	Y	Y	Y
Top 10 state (province) dummies	Y	Y	Y
U.S. dummy	Y	Y	Y
Mail week dummies	Y	Y	Y
Price elasticity	-0.28	0.71	0.24
Mean(dependent variable) for the Low Price group	0.155	530.9	82.6
Adj R-squared		0.645	0.031
N	1983	269	1983

Notes: This table reports the results (marginal effects) from the following regressions using observations from the pricing experiment. The dependent variables are (1) whether the case settles within 30 days; (2) the revenue for settled cases; and (3) the revenue of a case (coded as \$0 if not settled in 30 days). The independent variable is a dummy indicating the High Price group (the Low Price group is the omitted group). Baseline price is the licensing-fee portion of the requested amount. All models report robust standard errors, clustered by industry (defined by the two-digit SIC codes). ***, **, and * are, respectively, significance levels of 1%, 5%, and 10%.

marginal effect of High Price is -0.03 and marginally significant). However, the increase in the settlement probability is not large enough to offset the lower revenue for the cases that do settle. As a result, the expected 30-day revenue of an average case is \$13 lower for the Low Price group, but the difference is not statistically significant.²²

²²The expected revenue of a case is the settlement probability multiplied by the expected revenue, conditional on settling. We

Table 8 also shows that larger firms are significantly more likely to settle and less likely to negotiate a discount conditional on settling; thus, they are associated with a significantly higher expected revenue. Older firms are also significantly more likely to settle, but generate only a slightly higher revenue (p-value is 0.15). The coefficient for log(Baseline price) is not always significant, because it captures only residual differences in the licensing after controlling for the complete set of image-use dummies (e.g., taxes that differ by state).

Message experiment

Regression results of the message experiment are also consistent with the raw-data outcomes reported in table 6. The 30-day settlement probability of any of the four treatment groups is 12-19 percentage points higher than the control (column (1) in table 9), representing an increase of 85 to 135 percent. The incremental effect of including a deadline on the settlement rate, holding the rest of the message constant, is six percentage points, and the difference is statistically significant (p-value is 0.013, column (2) in table 9). For settled cases, the conditional settlement revenue is similar across all groups (columns (3) and (4) in table 9). In terms of the expected revenue per case (taking the settlement probability into account), the inclusion of a non-deadline message increases expected revenue by \$65, while a deadline message increases expected revenue by \$86 (columns (5) and (6) in table 9).

5.2 Results for different subsamples

Pricing experiment: comparing High Price and Low Price

Recall that, for an average case, a substantial reduction in the requested amount generates only a small change in the settlement probability. This result also holds across different subsamples. Figure 2a plots the marginal effects of the High Price dummy on the settlement probability by firm-size quintiles, with Q1 being the smallest firms in terms of annual sales and Q5 the largest (this and all other regression results for different subsamples are reported in the Online Appendix).²³ For all size quintiles, the settlement probabilities are not statistically different between the two price groups. Consequently, for all quintiles, the expected revenue is slightly higher for the High Price group, though none of the differences are statistically significant (not plotted here).

denote expected revenue as $R(p)$, where $R(p) = \Pr(p)\hat{R}(p)$; p is the requested amount; and $\hat{R}(p) \leq p$ because negotiation is allowed. The decomposition of the price elasticities is then $\frac{\partial R(p)}{\partial p} \frac{p}{R(p)} = \frac{\partial \Pr(p)}{\partial p} \frac{p}{\Pr(p)} + \frac{\partial \hat{R}(p)}{\partial p} \frac{p}{\hat{R}}$; that is, the price elasticity of the expected revenue is the sum of the elasticity of the settlement probability and the elasticity of the expected revenue conditional on settling. The two elasticities based on the regression results are -0.28 and 0.71, implying a positive elasticity of the expected revenue.

²³Firms at the 20th, 40th, 60th, and 80th percentiles have \$111K, \$190K, \$480K and \$1.7m in annual sales, respectively.

Table 9: Message effects on settlement outcomes

Dependent variable	Settle		Revenue for settled cases		Revenue	
	Logit		OLS		OLS	
	(1)	(2)	(3)	(4)	(5)	(6)
Waive the Cost	0.12*** (0.03)		-1.89 (17.50)		65.54*** (20.54)	
Waive the Cost + Deadline	0.18*** (0.03)		7.31 (17.05)		86.08*** (15.53)	
Discount	0.13*** (0.03)		17.11 (20.21)		66.67*** (16.93)	
Discount + Deadline	0.19*** (0.02)		3.72 (14.72)		88.54*** (15.60)	
Non-deadline message		0.12*** (0.03)		4.89 (17.21)		65.17*** (19.80)
Deadline message		0.18*** (0.03)		1.90 (16.99)		86.42*** (15.45)
Control variables	Y	Y	Y	Y	Y	Y
Mean(dependent variable) for the control group	0.14	0.14	485.2	485.2	68.1	68.1
Adj R-squared			0.81	0.81	0.07	0.07
N	2291	2291	534	534	2295	2295

Notes: This table reports the results (marginal effects) from three sets of regressions using the message-experiment sample. The dependent variables are, respectively, whether the case settles in 30 days; the revenue for settled cases; and the revenue of a case (coded as \$0 if not settled in 30 days). In columns (1), (3) and (5), the independent variables are four dummies indicating the treatment messages (the default is the control group). In the other columns, Non-deadline message equals one if the case belongs to either of the two non-deadline message groups; and Deadline message equals one if the case belongs to either of the two deadline message groups. Columns (3) and (4) use cases that settle in 30 days. In all regressions, the control variables are log(Requested amount), image count, log(Annual sales), log(Age+1), missing-age dummy, dummies indicating the top ten industries, the top ten states (provinces), the U.S., and the mail date. All models report robust standard errors, clustered by industry. ***, **, and * are, respectively, significant levels of 1%, 5%, and 10%.


Figure 2b plots the marginal effects of the High Price dummy on the settlement probability by the price-change percentage.²⁴ The effect of the price change on settlement probability increases monotonically in magnitude with this percentage and reaches statistical significance ($p < 0.05$) when the percentage is medium or high (above 0.70). The magnitude of the effect is, respectively, three and four percentage points when the price-change percentage is medium and high. However, even in these ranges, the effect of the price change on expected revenue per case is statistically indistinguishable from zero.

Message experiment: comparing the non-deadline messages and the control

As shown previously, for an average case, including a non-deadline message increases the settlement probability by 12 percentage points relative to the control group. This effect is relatively flat across different

²⁴Due to the presence of six discrete baseline licensing fees, we categorize 30 percent of the cases as low, 40 percent as medium, and 30 percent as high. The separating thresholds are 0.70 and 0.90.


Figure 2: Price effects on settlement probability for different subsamples



Notes: These figures report marginal effects of the High Price dummy on 30-day settlement probability (the baseline is the Low Price group). Q1-Q5 indicate firm-size quintiles based on annual sales, with Q1 including the smallest firms and Q5 the largest. Price-change percentage is $(\$400 \times \text{number of images}) / \text{total licensing fee}$. Bands indicate 95 percent confidence intervals.

subsamples, across both firm size and the relative size of the forgiven amount. Figure 3a shows that, relative to the control, adding a non-deadline message has an effect of at least seven percentage points for all firms, except those in the third quintile of the firm-size distribution. Compared to firms in the first quintile, the message effect for larger firms is not statistically different, except for firms in the top quintile, where it is 14 percentage points larger (p-value of the difference in the effect is 0.008).

Figure 3: Effects of a non-deadline message on settlement probability for different subsamples



Notes: This figure plots the marginal effect of a dummy variable that equals one if the case belongs to either of the two non-deadline message groups (the baseline is the control group) on the 30-day settlement probability. Q1 are the smallest firms based on annual sales and Q5 the largest. $\frac{\text{forgiven amount}}{\text{requested amount}} = (\$400 \times \text{number of images}) / \text{total licensing fee}$. Reported bands reflect 95% confidence intervals. The logit regressions (reported in table A3 in the Online Appendix) use cases in the control and the two non-deadline message groups. They include the interaction terms between the 'non-deadline message' dummy and firm-size quintile indicators (or indicators for different levels of the relative size of the forgiven amount).

Figure 3b shows that the effect of a non-deadline message is large and statistically significant, regardless

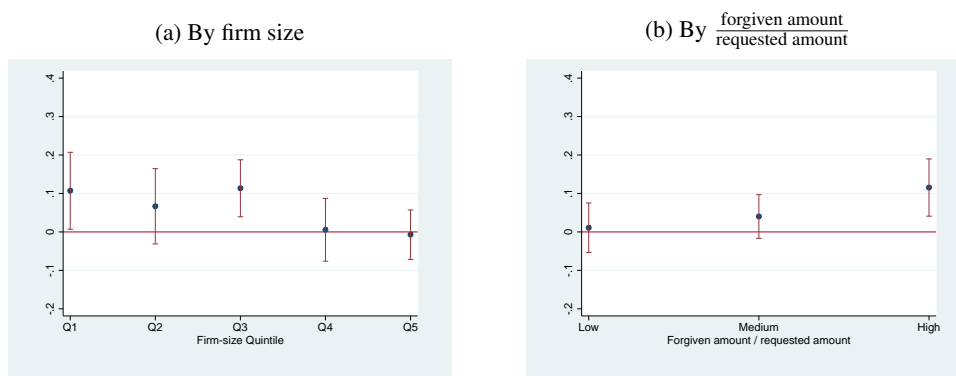
of the proportion of the forgiven amount relative to the requested amount. The differences in the effects between different ranges are small and not statistically different from zero.

Message experiment: comparing messages with and without a deadline

We evaluate the incremental impact of adding a deadline after which the price-reduction offer expires, holding the rest of the message fixed, on pre-deadline outcomes.²⁵ The deadline is set for the 22nd day after the first letter is mailed. In practice, the Agency did not take action immediately after the deadline; nor did it increase the requested amount in subsequent letters. Thus, we focus on the 22-day outcome to provide the cleanest comparison.

For an average case, the addition of a deadline increases the 22-day settlement likelihood by 5.4 percentage points (the p-value is 0.007). This effect, however, differs significantly among subsamples. The effect of a deadline on pre-deadline settlement probability is positive (roughly ten percentage points) for smaller firms in the lowest three quintiles, but not for larger firms in the top two quintiles (figure 4a). The effects of including a deadline for firms in the second and the third quintiles are statistically indistinguishable from those for firms in the first quintile, while effects in the fourth and fifth quintiles are significantly smaller (differences of ten and eleven percentage points, p-values 0.069 and 0.051).

Figure 4: Effect of a deadline on pre-deadline (22-day) settlement probability for different subsamples



Notes: This figure plots the marginal effects of a dummy that equals one if the case belongs to either of the two deadline message groups on the pre-deadline (22-day) settlement probability, holding the rest of the message constant. The baseline is comprised of cases belonging to either of the two non-deadline message groups. Q1 are the smallest firms based on annual sales and Q5 the largest. $\frac{\text{forgiven amount}}{\text{requested amount}} = (\$400 \times \text{number of images}) / \text{total licensing fee}$. Reported bands reflect 95% confidence intervals. The logit regressions (reported in table A4 in the Online Appendix) use cases from the four treatment groups of the message experiment. They include the interaction terms between the ‘deadline message’ dummy and firm-size quintile indicators (or indicators for different levels of the relative size of the forgiven amount).

²⁵Note that comparing outcomes for the groups with and without a deadline may underestimate the stand-alone effect of adding a deadline because infringers may already be motivated by the rest of the message to settle. One could alternatively create a second control group in which a deadline alone is added to the control letter.

Adding a deadline is significantly more effective when the potential future price increase is larger relative to the requested amount (figure 4b). For cases with the highest ratios, the pre-deadline settlement rate increases by more than ten percentage points (p-value of 0.0024) relative to the non-deadline messages; and the effect of a deadline is significantly different from that for cases with the lowest ratio (p-value of the difference is 0.002). To summarize, the addition of a deadline is not effective for the largest firms, and firms, on average, pay attention to the relative size of the potential price increase.

5.3 Plausible mechanisms

Recall that the model in section 4.1 captures two effects of pricing: the infringer's demand for settlement (the traditional price effect), and his belief about his potential future liability risk (the threat effect). Exogenous changes in (i) the actual price, and (ii) the provision of relevant price information, change the infringer's expected payoffs from complying or not. The contrast between the weak response to a substantial price reduction in the pricing experiment and the large and positive response to the messages (in particular, the sizable effect of including a deadline) is well-explained by the model. The lack of response to the price reduction happens because a lower request simultaneously sacrifices an implicit threat of escalation (that is, the traditional price effect and the threat effect offset each other). In contrast, the inclusion of a deadline appears to simultaneously strengthen the inducement of a lower price today and the signal of a willingness to escalate in the future. This conclusion is further supported by the heterogeneity in results for different sub-groups of infringers (see figures 2a and 4a together). While the price effect is small for all firms, adding a deadline has a significant effect for small firms, but little effect for larger firms. If larger firms are more price-inelastic than smaller firms and have more precise knowledge about legal risks, one expects them to be less influenced by a future price increase and its implied threat.

The effect of including a deadline on settlement rates also varies with the relative size of the forgiven amount (figure 4b). The effect is significantly greater when the forgiven amount is larger relative to the requested amount, which is consistent with a greater price inducement to settle sooner, and a larger increase in the potential future threat.

In contrast, the effect of the two non-deadline messages relative to the control group is pretty consistent across different subsamples (figure 3). Recall that the two messages contain two pieces of information that the control group lacks: (i) price information and (ii) acknowledgement of possible unintentionality. As discussed in Section 4.3, if price information is the main driver of the effect (through the expected-payoff mechanism), one may expect the effect of these messages to increase with firm size and the relative size of the forgiven amount. That is, the effect should exhibit heterogeneity similar to that observed when including a deadline. The lack of heterogeneity suggests that the expected-payoff mechanism alone is not sufficient

to explain this set of results. Together with the prevalence of unintentional infringement in this context, it is reasonable to conclude that the acknowledgement of possible unintentionality plays an important role in motivating infringers to settle.

5.4 Effects of the treatment messages on legal costs

Apart from increasing revenue for photographers, a higher settlement rate may also reduce the social costs of resolving disputes. If the infringer refuses to settle, the case is eventually sent to an external law firm for further collection. In this section, we provide a back-of-the-envelope calculation of the difference in legal costs associated with escalating a case to external enforcement that is attributable to the treatment messages. As very few cases end up in court, we do not consider the potential impact on litigation costs.

In the message experiment, over a 120-day horizon, 50.5 percent of the cases in the control group and 45.8 percent in the treatment groups are sent to external law firms, and the difference is statistically significant (p -value is 0.052). As explained previously, outcomes beyond 30 days may be affected by follow-up letters sent by the Agency. The actual difference in the probability of external enforcement is likely to significantly underestimate the impact of the messages because substantially more letters are sent for cases in the control group than in the treatment groups.²⁶ Given the available information, we use the actual difference of 4.6 percentage points as a lower-bound estimate and the difference in the 30-day settlement rate (12.4 percentage points) as an upper-bound estimate, assuming that the short-run effect is persistent. Thus, the messages reduce the probability of escalating to external enforcement by roughly 9.1-25 percent.

The main law firm with which the Agency contracts in the U.S. uses the following fee structure (reported with all details disguised for confidentiality). The Agency pays a fixed fee ($\$a$) for each case sent. The law firm typically adds an attorney fee of $\$c$ in the requested amount from infringers and deducts this amount from the collected revenue for a settled case. Furthermore, the Agency also pays a contingency fee of b percent of the remaining collected revenue for a settled case. Using this fee structure to approximate legal costs, we first calculate a settlement rate for cases sent to external law firms (d percent) based on historical data, consisting of infringement cases in 2011-2013, provided to us by the Agency.²⁷ Next, we assume that the average revenue (after deducting the attorney fee $\$c$) for settled cases is the same as for cases settled within 30 days of the message experiment ($\$505$). Finally, we assume that infringers hire legal counsel for

²⁶For example, 42 percent of the control group received a FedEx letter after about 30 days, while only 15 percent of the treatment groups did (table A9 in the Online Appendix). The settlement rate of the control group does catch up after 120 days. The timing of the start of the catchup (there is a discrete change in the slope of the cumulative settlement probability for the control group but not for the treatment groups) is after about 40 days, consistent with the firms responding to follow-up letters.

²⁷We use historical data for these statistics because enforcement through external law firms typically takes a long time, and the time frame of our experimental sample is not long enough to capture all associated outcomes.

20 percent of the cases, and that their attorney fee is also \$c per case.²⁸

Table 10: A back-of-the-envelope calculation of savings in legal costs

	Lower bound	Upper bound
<u>Panel 1: Assumptions</u>		
Number of cases per month in the U.S. and Canada	1,000	1,000
Difference in the number of cases sent to external law firms between the control and the treatment groups	46	124
Contractual terms by the Agency’s law firm (details disguised for confidentiality):		
Fixed fee per case sent		\$a
Attorney fee deducted from the collected revenue of a settled case		\$c
Contingency fee based on remaining revenue of a settled case (after deducting \$c)		b%
Settlement rate for cases sent to external law firms		d%
Average revenue per settled case after deducting the attorney fee \$c		505
Percentage of infringers obtaining legal service		0.2
Attorney fee per case charged by the infringers’ lawyers		\$c
<u>Panel 2: Results</u>		
Savings for the Agency (fixed and contingency fees)	\$866.5	\$2,335
Savings for the infringing firms	\$4,347	\$11,718
Total savings in legal costs	\$5,214	\$14,053

Notes: This table presents a calculation of savings in the legal costs, relative to the control group, due to the treatment messages for the population of cases in the message experiment. Take the upper-bound estimate, for example. The savings for the Agency consist of (1) fixed fees ($124 \times a$), and (2) contingency fees ($124 \times d\% \times 505 \times b\%$). The savings for the infringing firms consist of (1) the attorney fees for firms that employ legal counsel ($124 \times 20\% \times c$) and (2) attorney fees requested by the Agency’s law firm for firms that settle ($124 \times d\% \times c$). We observe the actual details of the legal fees and use them for all calculations in this table, but are not allowed to report them due to confidentiality.

For each 1,000 cases in the U.S. and Canada, the savings in legal costs are approximately \$866-\$2,335 for the Agency and \$4,347-\$11,718 for infringing firms (table 10). Total in-house enforcement revenue (revenue obtained without the use of external law firms) per thousand cases for the control group in the message experiment is \$192,000. Thus, savings in total legal costs are between 2.7 and 7.3 percent of the Agency’s in-house enforcement revenue. This may, in fact, underestimate the total savings in social costs due to additional factors, such as the opportunity cost of time and the psychological costs of stress for both parties from being entangled in prolonged disputes.

6 Conclusion

We study the effects of different enforcement methods on settlement outcomes for instances of copyright infringement by businesses. We use a novel, proprietary dataset generated from two large-scale field experiments run by a leading stock-photography agency. The experimental design avoids the confounding effects

²⁸We do not have information about the percentage of infringers that employ legal counsel after the Agency moves a case to an external law firm. The correspondence records between the Agency and firms that made contact show that, at the in-house enforcement stage, infringing firms are already employing external legal counsel eight percent of the time. Once the Agency escalates the case to outside law firms, we expect the percentage to be higher.

often encountered during real price changes, and we find that, on average, a substantial reduction in the requested amount alone generates only a small increase in the settlement probability. In contrast, for the same reduced request, the addition of a message that informs infringers of the price reduction and acknowledges the possibility of unintentional infringement has a large positive effect on settlement. Including a deadline further improves the settlement rate. Beyond generating higher revenues for copyright owners, the higher settlement rate resulting from the messages also substantially reduces the social costs of dispute resolution.

Typically, one would expect a large concession in price to be effective in resolving disputes, and pricing is often the sole focus of many bargaining models. Our findings suggest that pricing may affect bargaining outcomes in non-obvious ways, because it can reveal information about important factors that influence a party's outside option (e.g., in our particular context, the copyright enforcer's willingness to escalate the dispute). This signaling role of pricing is similar to findings in other situations with incomplete information (e.g., prices as signals for product quality, Wolinsky (1983)).

Our results also suggest that, when disputes are likely to be the result of inadvertent mistakes, proactive acknowledgment of this possibility may encourage collaborative outcomes (at least when the monetary stakes are small). This result may be generalized to other, similar contexts. Given the representativeness of the firms in our data, the effects that we document for digital images are likely to generalize to other types of copyrighted materials used by businesses, such as videos, music, and designs. With regard to consumer piracy, many contexts studied by the existing literature typically involve willful infringement (e.g., repeat file-sharers of music and movies). There are, however, areas in which potential infringers are unfamiliar with their legal obligations. An interesting example is the increasing use of 3D printing: many people download designs from the Internet without being aware of its illegality (Depoorter (2014)). Beyond copyright, inadvertent mistakes due to forgetfulness, disorganization, and limited awareness could also result in conflicts in a wide variety of contexts that may be mitigable using such strategies.

More generally, our results highlight the fact that communication can have an important impact on economic outcomes in some situations. Existing evidence from the lab and field are based primarily on individuals (see, e.g., Roth and Murnighan (1982) on the role of information in bargaining games; and DellaVigna and Gentzkow (2010) for a survey of studies on persuasive communication in a variety of contexts). Our study shows that these interventions can also influence firms. In this sense, it provides a small but important step towards deepening our understanding of these types of mechanisms beyond individual consumers in the broader economy.

References

- Alm, J. and B. Torgler (2011). Do ethics matter? tax compliance and morality. *Journal of Business Ethics* 101, 635–651.
- Anderson, E. and D. Simester (1998). The role of sale signs. *Marketing Science* 17, 139–155.
- Bebchuk, L. (1984). Litigation and settlement under imperfect information. *RAND Journal of Economics* 15, 404–415.
- Bessen, J. E., J. Ford, and M. J. Meurer (2011). The private and social costs of patent trolls. *Regulation* 34, 26–35.
- Bessen, J. E. and M. J. Meurer (2006, May). Patent Litigation with Endogenous Disputes. *American Economic Review* 96(2), 77–81.
- Bhattacharjee, S., R. D. Gopal, K. Lertwachara, and J. Marsden (2006). Impact of legal threats on individual behavior: An analysis of music industry actions and online music sharing. *Journal of Law and Economics* 49, 91–114.
- Cadena, X. and A. Schoar (2011). Remembering to pay? reminders vs. financial incentives for loan payments. Working paper.
- Cohen, L., U. Gurun, and S. D. Kominers (2014, August). Patent trolls: Evidence from targeted firms. HBS Working Paper 15-002.
- Danaher, B., M. Smith, R. Telang, and S. Chen (2014). The effect of graduated response anti-piracy laws on music sales: Evidence from a natural experiment in france. *Journal of Industrial Economics* 62, 541–553. forthcoming in *Journal of Industrial Economics*.
- Danzon, P. M. and L. A. Lillard (1983). Settlement out of court: The disposition of medical malpractice claims. *Journal of Legal Studies* 12, 345–377.
- DeBriyn, J. (2012). Shedding light on copyright trolls: An analysis of mass copyright litigation in the age of statutory damages. *UCLA Entertainment Law Review* 19.
- DellaVigna, S. and M. Gentzkow (2010). Persuasion: Empirical evidence. *Annual Review of Economics* 2, 643–669.
- Depoorter, B. (2014). Intellectual property infringements & 3d printing: Decentralized piracy. *Hastings Law Journal* 65, 1483–1503.
- Fehr, E., G. Kirchsteiger, and A. Riedl (1993, May). Does Fairness Prevent Market Clearing? An Experimental Investigation. *Quarterly Journal of Economics* 108(2), 437–59.
- Feld, L. P. and B. S. Frey (2002a). The tax authority and the taxpayer: An exploratory analysis.
- Feld, L. P. and B. S. Frey (2002b). Trust breeds trust: How taxpayers are treated. *Economics of Governance* 3, 87–99.
- Fellner, G., R. Sausgruber, and C. Traxler (2013). Testing enforcement strategies in the field: Threat, moral appeal and social information. *Journal of the European Economic Association* 11, 634–660.

- Fournier, G. M. and T. W. Zuehlke (1996). The timing of out-of-court settlements. *Rand Journal of Economics* 27, 310–321.
- Frey, B. S. (1997). *Not Just for the Money: An Economic Theory of Personal Motivation*. Cheltenham, UK and Brookfield, USA: Edward Elgar.
- Galasso, A., M. Schankerman, and C. Serrano (2013). Trading and enforcing patent rights. *RAND Journal of Economics* 44, 275–312.
- Glückler, J. and R. Panitz (2013). Survey of the global stock image market 2012. part i: Players, products, business. Technical report, Heidelberg: GSIM Research Group.
- Güth, W., R. Schmittberger, and B. Schwarze (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3, 367–388.
- Hallsworth, M., J. List, R. Metcalfe, and I. Vlaev (2014). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. Working Paper 20007, National Bureau of Economic Research.
- Inman, J. J., L. McAlister, and W. D. Hoyer (1990). Promotion signal: proxy for a price cut? *Journal of Consumer Research* 17, 74–81.
- Jolls, C., C. R. Sunstein, , and R. Thaler (1998). A behavioral approach to law and economics. *Stanford Law Review* 50, 1471–1550.
- Kahneman, D., J. L. Knetsch, and R. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review* 76, 728–741.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–291.
- Karlan, D., M. Morten, and J. Zinman (2015). A personal touch: Text messaging for loan repayment. NBER Working Paper No. 17952.
- Korobkin, R. and C. Guthrie (1994). Psychological barriers to litigation settlement: An experimental approach. *Michigan Law Review* 93, 107–192.
- Lanjouw, J. and M. Schankerman (2001). Characteristics of patent litigation: A window on competition. *RAND Journal of Economics* 32, 129–151.
- Lanjouw, J. O. and M. Schankerman (2004). Protecting intellectual property rights: are small firms handicapped? *Journal of Law and Economics* 47, 45–74.
- Lantagne, S. (2004). The morality of mp3s: The failure of the recording industry’s plan of attack. *Harvard Journal of Law & Technology* 18, 269–293.
- Luo, H. and J. H. Mortimer (2016, May). Copyright infringement in the market for digital images. *American Economic Review: Papers and Proceedings* 106(5), 140–145.
- Luttmer, E. F. P. and M. Singhal (2014). Tax morale. *Journal of Economic Perspectives* 28, 149–68.
- Meurer, M. (1989). The settlement of patent litigation. *RAND Journal of Economics* 20, 77–91.
- Perez-Truglia, R. and U. Troiano (2015, March). Tax debt enforcement: Theory and evidence from a field experiment in the united states.

- Priest, G. and B. Klein (1984). The selection of disputes for litigation. *Journal of Legal Studies* 13, 1–55.
- Rachlinski, J. J. (1996). Gains, losses, and the psychology of litigation. *Southern California Law Review* 70, 113–185.
- Reimers, I. (2015, February). Can private copyright protection be effective? evidence from book publishing.
- Roth, A. E. (1995). Bargaining experiments. In J. H. Kagel and A. E. Roth (Eds.), *Handbook of Experimental Economics*. Princeton: Princeton University Press.
- Roth, A. E. and J. K. Murnighan (1982, September). The Role of Information in Bargaining: An Experimental Study. *Econometrica* 50(5), 1123–42.
- Scott Morton, F. and C. Shapiro (2014). Strategic patent acquisitions. *Antitrust Law Journal* 79, 463–499.
- Spier, K. E. (2007). Litigation. In A. M. Polinsky and S. Shavell (Eds.), *Handbook of Law and Economics*, Chapter 4. North-Holland.
- Thaler, R. H. and C. R. Sunstein (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven & London: Yale University Press.
- Trubek, D. M., A. Sarat, W. L. Felstiner, H. M. Kritzer, and J. B. Grossman (1983). The costs of ordinary litigation. *UCLA Law Review* 31, 72–127.
- Waldfogel, J. (1995). The selection hypothesis and the relationship between trial and plaintiff victory. *Journal of Political Economy* 103, 229–260.
- Waldfogel, J. (2012). Copyright protection, technological change, and the quality of new products: Evidence from recorded music since napster. *Journal of Law & Economics* 55, 715–40.
- Wolinsky, A. (1983). Prices as signals of product quality. *Review of Economic Studies* 50, 647–658.

Online Appendix

Appendix A. Additional Tables and Figures

Table A1: Regression results on the settlement probability with and without using controls

	Logit model: DV = Settle			
	Pricing experiment		Message experiment	
	(1)	(2)	(3)	(4)
High Price dummy	-0.03*	-0.03*		
	(0.02)	(0.02)		
Waive the Cost			0.12***	0.11***
			(0.03)	(0.03)
Waive the Cost + Deadline			0.18***	0.16***
			(0.03)	(0.03)
Discount			0.13***	0.11***
			(0.03)	(0.03)
Discount + Deadline			0.19***	0.17***
			(0.02)	(0.02)
Control variables	Y	N	Y	N
Adj R-squared	0.06	0.002	0.06	0.02
N	1983	1983	2291	2295

Notes: This table reports the logit regression results on the 30-day settlement probability without using any controls. For comparison, the regression results using controls are replicated in the first and the third columns (that is, the results reported in column (1) in tables 8 and 9). All models report robust standard errors, clustered by industry. ***, **, and * are, respectively, significant at levels of 1%, 5%, and 10%.

Table A2: Price effects on settlement outcomes for different subsamples

Dependent variable	Settle		Revenue for settled cases		Revenue	
	Logit		OLS		OLS	
	(1)	(2)	(3)	(4)	(5)	(6)
High Price dummy	-0.02 (0.04)	-0.00 (0.04)	155.77 (98.81)	324.20*** (52.12)	8.77 (14.34)	19.63 (27.13)
Size quintile 2 × High Price		-0.00 (0.06)	30.39 (111.78)		3.11 (28.35)	
Size quintile 3 × High Price		-0.02 (0.06)	188.53 (115.43)		11.08 (26.16)	
Size quintile 4 × High Price		0.00 (0.04)	66.22 (115.06)		16.55 (22.53)	
Size quintile 5 × High Price		-0.02 (0.05)	170.03 (106.42)		-0.80 (45.26)	
Size quintile 2		0.01 (0.06)	80.13 (91.75)		13.33 (31.68)	
Size quintile 3		0.04 (0.06)	7.06 (83.24)		8.78 (24.63)	
Size quintile 4		0.04 (0.05)	108.86 (90.01)		41.95 (32.48)	
Size quintile 5		0.06 (0.08)	130.29 (123.70)		92.71 (60.51)	
Medium price-change % × High Price		-0.03 (0.04)		-89.41 (63.42)		-11.95 (29.47)
High price-change % × High Price		-0.04 (0.04)		-92.44 (63.01)		-0.66 (32.65)
Medium price-change %		0.05 (0.05)		-104.84 (164.30)		-1.90 (61.53)
High price-change %		0.13* (0.07)		-27.85 (245.03)		57.20 (102.86)
Control variables	Y	Y	Y	Y	Y	Y
Adj R-squared			0.653	0.646	0.030	0.031
N	1983	1983	269	269	1983	1983

Notes: This table reports the marginal effects of three sets of regressions using the observations in the pricing experiment, in which the dependent variables are (1) whether the case settles in 30 days; (2) the revenue for settled cases; and (3) the revenue of a case (coded as \$0 if not settled in 30 days). In all regressions, the independent variable is a dummy indicating the High Price group (the Low Price group is the default). Size quintiles are based on annual sales, with quintile 1 including the smallest firms. Price-change percentage is defined by $(\$400 \times \text{number of images}) / \text{licensing fee}$. The controls in all regressions include $\log(\text{Baseline price})$, image count, $\log(\text{Annual sales})$, $\log(\text{Age}+1)$, a missing-age dummy, and dummies indicating the top ten industries (defined by the two-digit SIC codes), the top ten states (provinces), a U.S. location, and the mail week. All models report robust standard errors, clustered by industry. ***, **, and * are, respectively, significant at levels of 1%, 5%, and 10%.

Table A3: Comparing the two non-deadline messages and the control for different subsamples

	Logit model: DV = Settle		
	(1)	(2)	(3)
Non-deadline message	0.10*** (0.02)	0.07 (0.05)	0.11*** (0.04)
Size quintile 2 × Non-deadline message		0.06 (0.09)	
Size quintile 3 × Non-deadline message		-0.08 (0.08)	
Size quintile 4 × Non-deadline message		0.00 (0.07)	
Size quintile 5 × Non-deadline message		0.14*** (0.05)	
Size quintile 2		-0.04 (0.06)	
Size quintile 3		0.08 (0.05)	
Size quintile 4		0.05 (0.05)	
Size quintile 5		0.00 (0.07)	
Medium ratio × Non-deadline message			0.02 (0.05)
High ratio × Non-deadline message			-0.04 (0.04)
Medium ratio			-0.03 (0.06)
High ratio			0.01 (0.08)
Control variables	Y	Y	Y
N	1420	1420	1420

Notes: This table reports logit regression results (marginal effects) using cases in the two non-deadline messages and the control groups in the message experiment. The dependent variable is whether the case settles in 30 days. Non-deadline message equals one if the case belongs to either of the two non-deadline message groups (the default is the control group). Size quintiles are based on annual sales, with quintile 1 including the smallest firms. Medium ratio equals one if $0.7 < \frac{\text{forgiven amount}}{\text{requested amount}} < 0.9$; and high ratio equals one if the ratio is higher than 0.9. In all regressions, the control variables are $\log(\text{Requested amount})$, image count, $\log(\text{Annual sales})$, $\log(\text{Age}+1)$, missing-age dummy, dummies indicating the top ten industries (defined by the two-digit SIC codes), the top ten states (provinces), the U.S., and the mail date. All models report robust standard errors, clustered by industry. ***, **, and * are, respectively, significant at levels of 1%, 5%, and 10%.

Table A4: Comparing messages with and without a deadline for different subsamples

	Logit model: DV = Settle (in 22 days)		
	(1)	(2)	(3)
Deadline message	0.05*** (0.02)	0.11** (0.05)	0.01 (0.03)
Size quintile 2 × Deadline message		-0.04 (0.06)	
Size quintile 3 × Deadline message		0.01 (0.06)	
Size quintile 4 × Deadline message		-0.10* (0.06)	
Size quintile 5 × Deadline message		-0.11* (0.06)	
Size quintile 2		0.08 (0.06)	
Size quintile 3		0.05 (0.05)	
Size quintile 4		0.13** (0.05)	
Size quintile 5		0.20*** (0.07)	
Medium ratio × Deadline message			0.03 (0.05)
High ratio × Deadline message			0.10*** (0.03)
Medium ratio			0.02 (0.05)
High ratio			0.01 (0.05)
Control variables	Y	Y	Y
N	1696	1696	1696

Notes: This table reports logit regression results (marginal effects) using cases in all four treatment groups in the message experiment. The dependent variable is whether the case settles in the first 22 days. Deadline message equals one if the case belongs to either of the two deadline message groups (the default group are cases that belong to either of the two non-deadline messages). Size quintiles are based on annual sales, with quintile 1 including the smallest firms. Medium ratio equals one if $0.7 < \frac{\text{forgiven amount}}{\text{requested amount}} < 0.9$; and high ratio equals one if the ratio is higher than 0.9. In all regressions, the control variables are a dummy indicating the discount framing, $\log(\text{Requested amount})$, image count, $\log(\text{Annual sales})$, $\log(\text{Age}+1)$, missing-age dummy, dummies indicating the top ten industries (defined by the two-digit SIC codes), the top ten states (provinces), the U.S., and the mail date. All models report robust standard errors, clustered by industry. ***, **, and * are, respectively, significant at levels of 1%, 5%, and 10%.

Table A5: Balance check for the ‘unfiltered’ sample

(a) Pricing experiment

Group	N	Baseline price	Sales				Age			
			Sales, \$1,000		log(Sales, \$)	Age		log(Age+ 1)		
			N	mean	median	mean	N	mean	median	mean
Low Price	945	662.3	896	25749	330	13.1	910	19.5	11	2.5
High Price	1,800	658.6 (0.85)	1,649	13020 (0.20)	290 (0.41)	13.0 (0.15)	1,593	18.1 (0.16)	11 (0.71)	2.5 (0.70)

(b) Message experiment

Group (p-value, diff. to control)	N	Baseline price (= Requested amount)	Sales				Age			
			Sales, \$1,000		log(Sales, \$)	Age		log(Age+ 1)		
			N	mean	median	mean	N	mean	median	mean
Control	866	639.3	754	25420	383	13.3	756	19.2	11	2.6
Waive the Cost	549	675.3 (0.15)	458	44562 (0.32)	348 (0.64)	13.4 (0.86)	455	19.5 (0.82)	11 (0.88)	2.6 (0.91)
Waive the Cost + Deadline	564	612.0 (0.24)	493	11023 (0.28)	300 (0.13)	13.1 (0.07)	494	16.9 (0.08)	9 (0.16)	2.5 (0.09)
Discount	571	609.7 (0.19)	526	10071 (0.23)	320 (0.26)	13.2 (0.21)	528	18.2 (0.42)	12 (0.80)	2.5 (0.75)
Discount + Deadline	577	652.6 (0.79)	526	48841 (0.55)	378 (0.89)	13.1 (0.01)	533	17.6 (0.19)	10 (0.30)	2.5 (0.32)

Notes: This table reports the balance tests using the ‘unfiltered’ sample (that is, without dropping cases for which the mail is returned or sales information is missing, etc). Note, however, that tier-4 and tier-5 firms in the pricing experiment are excluded to ensure that firms in the two experiments are comparable in size. The test statistics are consistent with those reported in the paper after dropping some types of cases.

Table A6: 30-day settlement outcomes for the ‘unfiltered’ sample

(a) Pricing experiment

	Low Price	High Price	Difference	p-value
Settlement probability	0.12	0.10	-0.02	(0.06)
Revenue for settled cases	528.09	1080.1	552.1	(0.02)
Revenue per case	65.94	109.8	43.9	(0.12)
N	945	1,800		

(b) Message experiment

Group	N	Settlement probability	Difference (compared to control)	p-value
Control	866	0.12		
Waive the Cost	549	0.19	0.07	(0.00)
Waive the Cost + Deadline	564	0.25	0.13	(0.00)
Discount	571	0.21	0.09	(0.00)
Discount + Deadline	577	0.27	0.15	(0.00)

Notes: This table reports the settlement outcomes for different groups in the two experiments using the ‘unfiltered’ sample (defined in table A5). The results are consistent with those reported in the paper after dropping some types of cases.

Appendix B: The stock photography industry

The stock photography industry provides images that are readily available for licensing. On behalf of individual photographers or organizations (such as National Geographic), stock photography agencies manage, market and license images to business customers, such as advertising and graphic design agencies, publishers and corporations.

Rights-managed (RM) and royalty-free (RF) are two typical license types. An RM license allows the licensee a one-time use of the image (e.g., 100 billboards in Germany for three months). For a given image, the licensing fee depends on the specified use. A licensee who wants to use the image beyond the specified scope must purchase a new license. In contrast, an RF license gives the licensee the right to use the image with few restrictions after making a one-time payment. Under the RF model, some agencies price images on an à la carte basis, while others use subscription pricing. According to a survey by Graphic Design USA, 32, 98, and 54 percent of responding designers had licensed images under the RM, à la carte RF, and subscription RF models, respectively, in 2011.

Stock images are categorized as either editorial or creative, depending on whether commercial use is allowed. Editorial images portray specific people, places, things and events that provide context for newspaper and magazine articles, blog posts, and other non-commercial presentations. Editorial images cannot be licensed for advertising or promotional purposes because there are typically no permissions from the subjects of these photos. In contrast, creative images are produced and staged by the photographer. Creative images can be licensed for both non-commercial and commercial purposes because they feature generic objects, or subjects who have consented to commercial use.

The industry is divided into traditional (premium) and micro-stock segments. Leading agencies in the traditional segment include Corbis Images and Getty Images; leading agencies in the micro-stock segment are Fotolia, iStockphoto and Shutterstock. Suppliers in the traditional segment are typically professional photographers who target customers and projects that require high-quality images. Images in this segment are licensed under the RM or the à-la-carte RF model. Depending on the quality of the image and the specified use, the licensing fee ranges from a couple of hundred to tens of thousands of dollars.²⁹ In addition to facilitating transactions, agencies in this segment also provide complementary services, including searching for images with specific themes, large-project management, rights clearance and protection against copy-right liabilities.

Agencies in the micro-stock segment source their images from a wide range of photographers, including professionals and hobbyists. Images in this segment are licensed under the RF model. The price level is significantly lower than in the traditional segment. For example, the vast majority of images on iStockphoto (à la carte pricing) are priced between \$2 and \$100 per image, depending on the size of the digital file; and, at Shutterstock (subscription pricing), the user can download up to 25 images per day for a monthly fee of \$249.

²⁹The most expensive cases usually occur because the licensed use is exclusive.

Appendix C. Follow-up letters

The Agency sends three types of follow-up letters. First, a second settlement request letter is sent after about 30 days unless the case has been settled by then. Second, also after 30 days or so, an extra FedEx letter is sent to a set of relatively large firms that have not yet responded. (The first and the second letters are sent by USPS standard mail). Third, after 60 days, the Agency sends a final notice by email if the firm has ignored the previous letters. Cases that do not settle (either because the firms ignore all the letters or do not settle after making contact) are eventually sent to external law firms. In our sample, about half of all cases are sent to external law firms, with 90 percent sent between 58 and 159 days after the initial letter.

The proportion of cases receiving the second letter is determined by the initial settlement outcomes of a particular group of infringers. Normally, the extra FedEx letter and the last email reminder are sent only if the infringer ignores the settlement request letter. In these cases, a ‘case-handler’ field is coded as empty. During the trial period, the compliance team at the Agency records the group to which a case belongs also in the ‘case-handler’ field. However, this group information is *not* recorded for the group designated as the ‘default’ group in each of the experiments (i.e., the High Price group in the pricing experiment and the control group in the message experiment). Thus, during the trial periods, the non-default groups do not get these two types of follow-up letters because the ‘case-handler’ field is not empty, even for infringers that have not made contact. This is sometimes corrected, but not always.

Table A7 reports, by group, the percentage of cases that receive the three types of follow-up letters in the two experiments. The differences suggest that the potential upward bias for the ‘default’ group is likely to be substantial for the message experiment. These incidences reflect the complexity and constraints that firms face when conducting field experiments. It is, at least, fortunate that we know the exact data-generating process, so we can discover any confounding factors.

Table A7: Follow-up letters by group

	Pricing experiment		Message experiment				
	High Price	Low Price	Control	Waive the Cost	Waive the Cost + Deadline	Discount	Discount + Deadline
Second letter	82.6%	79.6%	82.8%	72.8%	66.7%	71.7%	64.3%
FedEx letter	5.2%	0%	41.9%	16.3%	15.6%	14.5%	14.2%
Final email	7.9%	0%	14.1%	0%	0%	0%	0%

Notes: This table reports, by group, the proportion of cases that receive different types of follow-up letters.

Appendix D. 120-day outcomes

This section reports the 120-day outcomes for both experiments. As explained in the paper, settlement outcomes after 30 days are confounded by different intensities of follow-up letters for different groups. Table A7 in Appendix C shows that the difference in the intensity of follow-up letters is relatively small between the two groups in the pricing experiment and, thus, we evaluate the 120-day performance using the actual data without further adjustment. Table A8a shows that removing the \$400 per-image surcharge results in a seven-percentage-points increase in the 120-day settlement probability. The increase in the settlement probability, however, is not enough to outweigh the decrease in the revenue conditional on settlement (\$211). Thus, consistent with the 30-day results, the expected revenue per case is statistically similar between the two groups (the lower-priced group generates \$11 lower, but the p-value of the difference is 0.49).

Table A8: Price and message effects on the 120-day outcomes

(a) Pricing experiment				
	Low Price group	High Price group	Difference	(p-value)
Settlement probability	0.32	0.25	-0.07	(0.00)
Revenue for settled cases	530.9	741.5	210.6	(0.00)
Revenue per case	167.9	178.9	11.0	(0.49)

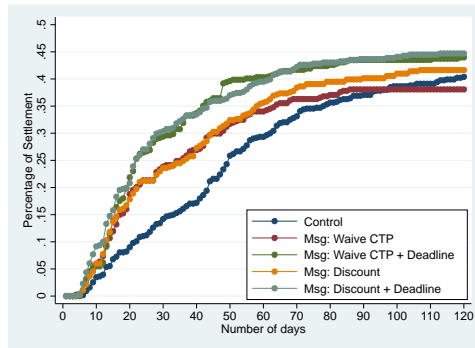
(b) Message experiment				
(p-value is relative to the control)	120-day settlement probability	120-day difference (lower bound)	30-day difference (upper bound)	
Control	0.40			
Non-deadline messages	0.40	0.00 (0.76)	0.10	(0.00)
Deadline messages	0.44	0.03 (0.20)	0.15	(0.00)

Notes: This table reports the 120-day settlement outcomes for different groups in both experiments. For the message experiment, we report only the effects on the settlement probability because the requested amount is similar across groups given the same usage and, thus, the effects on the expected revenue per case are consistent with that on the settlement probability.

The differences in the follow-up intensity between the control and the treatment groups, however, are substantial in the message experiment. Table A7 in Appendix C shows that 82 percent of the control group received the second letter, while 64-72 percent in the treatment groups did; 42 percent of the control group received a FedEx letter, while only 15 percent of the treatment groups did; and 14 percent of the control group received the final email reminder, while none of the cases in the four treatment groups did. According to the Agency, the FedEx letters are, perhaps, the most powerful follow-up action because (1) they go to relatively large firms that are more likely to eventually settle; and (2) they are hand-signed by the head of the compliance team and sent through FedEx express, while, in contrast, all other letters are sent either by standard USPS mail or via email and without the signature. Figure A1 plots the cumulative settlement probability in 120 days for the five groups in the message experiment. The results show that the control group catches up after 30 days. The timing of the start of the catchup (i.e., the discrete change in the slope for the control group) is after about 40 days, which is consistent with the firms responding to follow-up letters. Therefore, the actual difference in the 120-day settlement rate is likely to substantially under-estimate the

effect of these messages.

Figure A1: Cumulative settlement probability by message (120-day)



Given the available information, we provide a bound estimate for the 120-day effect of the messages, using the actual difference as the lower bound and the 30-day difference as the upper bound, assuming that the short-run effect is persistent. Table A8b shows that, in the message experiment, the two messages that do not impose a deadline perform similarly to the control group after 120 days, while the two messages that impose a deadline yield a settlement rate that is three percentage points higher than the control (p-value is 0.20). Combining these actual differences with the 30-day estimates, the 120-day effect of the non-deadline messages is estimated to be between 0 and 10 percentage points (0-25 percent increase relative to the control group), and that of the two deadline messages is 3 to 15 percentage points (an increase of 7.5-33 percent relative to the control).