Information Design in Operations Management

HARVARD UNIVERSITY
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the Committee for the PhD in Business Administration have examined a dissertation entitled

Information Design in Operations Management

Presented by MoonSoo Choi

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature

Ryan Buell, Chair, Dissertation

Signature

Leslie John

Signature

Antonio Moreno

Date September 2, 2021

Dissertation Advisor:
Ryan Buell

Author:
MoonSoo Choi

Information Design in Operations Management

ABSTRACT

Information design – or the practice of effectively communicating information to its audience – is a delicate dance that requires accuracy, clarity, comprehensiveness, and engagement. Yet, even when all these boxes are checked off, it is no secret that people process information differently. The different ways people interpret and assimilate information, in turn, lead to a wide, non-uniform range of choices and decisions.

Such is also the case in firm operations. Many firms serve a broad spectrum of customers with different needs, preferences, and behaviors. As a result, although it is crucial for firms to examine how different types of information design affect consumer perceptions and behavior, customer heterogeneity often hinders firms from measuring the impact of information design at a holistic, systemic, and operational level. This dissertation seeks to tackle such research question; specifically, the dissertation examines how different information design choices affect consumer perceptions, consumer engagement, and firm operations. The dissertation leverages methods and techniques in field experiments, causal inference, machine learning, and lab experiments.

The dissertation consists of three distinct essays. The first chapter explores how providing transparency into tradeoffs affects customer acquisition, engagement, and retention, through running a field experiment with 393,036 customers considering applying for bank-issued credit cards. The results reveal that – although providing transparency into tradeoffs did not have

a significant effect on customer acquisition rates – customers who were exposed to the transparency exhibited increased product usage, higher retention, and lower likelihood of making late payments. In the presence of a promotional campaign that provides financial incentives, however, the positive effects of tradeoff transparency were attenuated. Moreover, the findings suggest that the effect of tradeoff transparency on consumer spending and retention is more pronounced among the customers with more familiarity and experience with credit card products. Overall, the results show that disclosing tradeoffs can be an effective strategy for firms to keep the customers better informed and improve customer engagement.

The second chapter examines how different facets of eCommerce delivery data can be used to develop and improve delay prediction models. We apply a mix of causal inference and machine learning (e.g., random forest) models to a comprehensive, large-scale dataset spanning user, delivery, and order information from JD.com, one of the largest eCommerce companies in China. In doing so, we first analyze how duration of each leg of the delivery, time allotted for each leg, and probability of delay relate to each other. Then, we fit random forest models to predict delays and identify primary predictors for such delays. Testing random forest models with different feature sets shows that including information about the earlier leg or warehouse package load can significantly improve the accuracy of the prediction model. Our prediction models suggest that managers can leverage various operational data to identify delays early on to prevent the orders from being delayed.

Finally, the third chapter seeks to tackle the perennial issue in consumer contracts: the 'no-reading problem.' Past studies have shown that somewhere between 74% and 99.8% of readers skip reading consumer contracts (or "fine print" or "terms and conditions"). In this chapter, we propose – and evaluate the effectiveness of – showing simpler translations of long, convoluted consumer contracts alongside the original version. Results of four experimental studies suggest that highlighting the key points of each clause of the contract significantly improves consumers' perceptions of the company, trust, and willingness to sign contracts. However, these effects were attenuated in the presence of risk associated with the contract and its associated service or product. We conclude the chapter with a general discussion on how managers and firms can best put our research findings into practice.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

Long yet short, dark yet enlightening, humbling yet empowering doctoral journey has come to a wrap – with countless people to thank from these past four years. I am immensely grateful to have been part of the HBS community and to have crossed paths with many dazzlingly smart yet immeasurably kind individuals.

To my committee: Ryan – you are an inspiration, period. It has been my absolute privilege to have had you as my advisor, mentor, and friend. Even when I was (and when others were) wrapped in the darkest of the shadows, you always knew how to cheer me and others up with unwavering encouragement, support, and empathy. I am grateful for having had the opportunity to bounce thoughts off with, seek advice from, and watch and learn from you. You are an exemplar in the community, and I will neither be the first nor the last to say this.

Toni – we both joined HBS the same year, and the timing could not have been more perfect. I have benefitted tremendously from our conversations and interactions: in RCTOM, Empirical TOM, COER, INFORMS, and other research workshops. Your research gave me first glimpse into research in retail operations, and I owe my thanks to you for my newly found interest in the topic. Your advice on research, writing, or life are golden, and I am thankful for your guidance throughout the years.

Leslie – your insights and wits are sharper than the sharpest cheddar. I owe you many thanks for your thoughts and feedback, which have been instrumental in development of my research. As I have been working on the last chapter of my dissertation, I have come across quite a few of your past works, which have immensely helped me think about information disclosure to consumers. Thank you!

To Friends of TOM: thank you. Ehsan, for your quirky quips and your brutal yet perceptive frankness, at Roksana, at Cotting, or across Arizona. Kala, for being an awesome biking & boba buddy and for being someone I can rely on. Ashley, for hosting numerous socials in your apartment, and for helping me realize the importance of taking a step back. Ryan & Rose, for our exchanges and for always welcoming me with warmth. Maya & Natalie, for venturing through thick and thin as operations scholars. Caleb, for being an amazing co-TF and for inspiring me to be a better researcher. Daniel, for being an active, vibrant TOM member within weeks of starting. Jeff, Ohchan, Tommy, Sourobh, Meitong, Grace, and James for helping me tackle problems from unseen perspectives. Michelle and Raha for sharing stories outside of research life. Justine, Paige, Omar, and Matt for giving us confidence that the future of TOM will be in good hands.

There are also many other members of the TOM unit whom I must thank. Kris, for your radiant energy, sharing your research in ML, and for letting me witness your uncanny ability to break down technical materials. Mike & RCTOM team, for allowing me to be RCTOM TF, witness the (countless hours of) inner workings of teaching preparation, and to learn to become a better teacher. Feng, for hosting lunch seminars and for ensuring that all TOM doctoral students can access and mobilize resources. Ananth, for your warmth and wisdom, and for inviting me to attend COER sessions. Matt Smith, Laura Marshman, Mandy Hagee, et al. for the undiminished positivity.

To the rest of the HBS research community. NERD Lab, for encouraging everyone to unleash maximal creativity. Hashim for being my thought partner throughout and for navigating data science together. Yosub, Do Yoon, Jaewon, Jon P., Olivia, Byungyeon, Jimin for exchanging research ideas, late-night sparks and conversations, and for philosophical musings. Jen, Marais, Angela, Kathy, Maryna, and LuAnn for your unconditional warmth, empathy, and understanding. Juana for your unfading smiles.

To many colleagues outside of the HBS community. Monika, for being an awesome friend, colleague, and co-author, and for our uplifting interactions, despite the 6-7 hours of time zone difference. Park, RJ, and Dayton – for bringing joy, especially in intermittently dull moments in conferences. BOMSI 2019 attendees – for sharing common interest in behavioral operations and for the fun, energizing summer session. Setareh – for helping me see the wisdom of life. Professor Ward Whitt – for introducing me to the world of research, and for encouraging me at my lowest points. Professor Rhonda Righter – for being my longtime mentor, and for introducing me to joy of studying probability. And, from the bottom of my heart, my sincerest thanks go out to my friends, family, colleagues, and many not mentioned here but are nonetheless close to my heart.

Ubi Caritas et Amor, Deus Ibi Est.

Improving Customer Compatibility with Tradeoff Transparency

Ryan W. Buell and MoonSoo Choi
Harvard Business School, Harvard University, Boston, Massachusetts 02163
[rbuell@hbs.edu, mochoi@hbs.edu]

Through a large-scale field experiment with 393,036 customers considering opening a credit card account with a nationwide retail bank, we investigate how providing transparency into an offering's tradeoffs affects subsequent rates of customer acquisition and long-run engagement. Although we find tradeoff transparency to have an insignificant effect on acquisition rates, customers who were shown each offering's tradeoffs selected different products than those who were not. Moreover, prospective customers who experienced transparency and subsequently chose to open an account went on to exhibit higher quality service relationships over time. Monthly spending was 9.9% higher and cancellation rates were 20.5% lower among those who experienced transparency into each offering's tradeoffs. Increased product usage and retention accrued disproportionately to customers with prior category experience – more experienced customers who were provided transparency spent 19.2% more on a monthly basis and were 33.7% less likely to defect after nine months. Importantly, we find that these gains in engagement and retention do not come at the expense of customers' financial wellbeing – the probability of making late payments was reduced among customers who experienced transparency. We further find that the positive effects of tradeoff transparency on engagement and retention were attenuated in the presence of a promotion that provided financial incentives to choose particular offerings. Taken together, these results suggest that providing transparency into an offering's tradeoffs may be an effective strategy for informing customer choices, leading to better outcomes for customers and firms alike.

[Keywords: Behavioral operations, operational transparency, customer compatibility, customer behavior]

1. Introduction

It is a common practice that when service companies market their offerings to prospective customers, they emphasize the advantages of each offering and downplay its tradeoffs. This strategy seems quite rational to the extent its objective is to convert browsers into buyers, as accentuating the downsides of an offering might drive some customers away. However, since value in many service settings arises from the long-term engagement and retention of satisfied customers, marketing messages that provide transparency into a service offering's tradeoffs (both its advantages *and its corresponding disadvantages*) may help customers make better-informed choices, facilitating the acquisition of more compatible customers, and improving long-run outcomes for customers and firms alike. Although this proposition may have intuitive appeal, the effect of providing prospective customers with transparency into an offering's tradeoffs has never been empirically investigated. In this paper, we introduce the concept of *tradeoff transparency*, which we define as voluntarily marketing the downsides of an offering with similar emphasis as its corresponding advantages, and we present the results of a large-scale field experiment in which we tested the impact of tradeoff transparency on customer acquisition and long-term engagement.

In service contexts, customers, who have heterogeneous needs and preferences (Frei 2006), and who provide key inputs to operating processes (Roels 2014, Sampson and Froehle 2006), are typically responsible for selecting the companies and service offerings with which they engage. This selection process is fraught with information asymmetry (Buell et al. 2016, Schmidt and Buell 2017), so much so that services have been classified as "experience goods," in that customers are rarely able to accurately assess their characteristics prospectively (Israel 2005). Yet service outcomes, like satisfaction and profitability, have been shown to hinge largely on customer compatibility – the degree of fit between the customer's individual needs and preferences, and the service offering's attributes (Buell, Campbell, et al. 2020). Hence, initiatives that reduce information asymmetry by making the tradeoffs of service offerings clearer to prospective customers may help those customers make better-informed decisions, leading to better long-term customer engagement. On the other hand, emphasizing an offering's tradeoffs may drive customers away, reducing customer acquisition, and undermining company performance.

To investigate this tension, we collaborated with Commonwealth Bank of Australia (CBA), the largest bank in the southern hemisphere (Buell and John 2018), and conducted a field experiment with 393,036 customers who were considering applying for bank-issued credit cards. Bank-issued credit cards are a relevant setting for this research because 1) credit cards vary in their design parameters, such that some offerings may be a better fit for an individual customer's needs and preferences than others, 2) credit card marketing typically emphasizes the advantages of each offering, and 3) customer engagement with credit cards (e.g., spend, late payments, retention, etc.) is a meaningful driver of long-term performance. Similar dynamics exist in a broad array of service contexts, where long-term outcomes are influenced by the engagement dynamics of heterogeneous customers. We note that since banks like CBA are legally required to provide customers with disclosure statements that reveal the terms and conditions of their offerings, this setting may be a conservative domain in which to study these effects. After all, tradeoff transparency in this context merely increases the salience of information about each offering's downsides, rather than presenting previously-undisclosed information. However, we note that even in settings where such disclosures are not mandated, online reviews that reveal the otherwise-hidden downsides experienced by prior purchasers are often available to customers. Hence, in most relevant contexts, as in this study, tradeoff transparency increases the salience of an offering's downsides.

In our field experiment, each prospective credit card customer was randomly assigned to either a control condition, in which they received traditional marketing messaging that emphasized the advantages of each credit card (e.g., our lowest interest rate on purchases, includes travel insurance, etc.), or a treatment condition, in which they received messaging that emphasized the same advantages, while also providing transparency into each card's tradeoffs (e.g., carries a higher annual fee, does not earn awards points, etc.). We subsequently tracked the progression of customers in both conditions through the acquisition funnel and documented the engagement of customers who chose to open accounts for nine months after activation, to assess the impact of tradeoff transparency on service performance.

Although we find that providing tradeoff transparency in this setting had no impact on overall rates of customer acquisition, we find that it had a meaningful effect on customers' choices, usage, and retention. Holding constant other factors, customers who experienced tradeoff transparency selected a different mix of cards, spent 9.9% more on a monthly basis, and after nine months, were 20.5% less likely to cancel their accounts. Importantly, this increased engagement did not come at the expense of customers' financial wellbeing – the probability of making late payments declined 10.8% in the initial six months of the service relationship among those who experienced tradeoff transparency. We further find evidence that the effects of tradeoff transparency on customer engagement are especially acute among customers with prior-category experience. Providing tradeoff transparency to below median age prospects, who had limited, if any, previous experience with credit cards, had a *de minimis* effect on usage and retention. In contrast, providing tradeoff transparency to older, more experienced customers, led to significantly higher levels of engagement – increasing monthly spending by 19.2% and decreasing the probability of cancellation after nine months by 33.7%. Finally, we observe that the effects of tradeoff transparency on subsequent usage and retention are attenuated among customers who were attracted by a promotion. Midway through the experimental period, the bank introduced a financial incentive for customers who opened specific types of credit cards, and customers who opened accounts during this promotion period exhibited a statistically indistinguishable response to the treatment. The results are consistent with the idea that customers who organically seek out an offering are more likely to consider and benefit from tradeoff transparency than those who are attracted or influenced by a promotion. By examining both the short-term marketing impact and the longer-run operational impact of providing prospective customers with tradeoff transparency, the present study lends support to prior research calling for an integrated approach to marketing and operations (Ho and Zheng 2004). Moreover, it builds on a growing literature that demonstrates how voluntarily revealing facets of an operation that are traditionally kept hidden may in some cases improve outcomes for customers and service providers alike (Buell et al. 2017, Mohan et al. 2020).

2. Tradeoff transparency, customer acquisition, and engagement

In this paper, we define tradeoff transparency as voluntarily marketing the downsides of an offering with similar emphasis as its corresponding advantages. As such, the concept of tradeoff transparency is closely related to prior research on different types of customer-facing transparency, information provision, and information salience, all of which have been closely studied in the operations and marketing literatures. In these contexts, transparency can be characterized as operating in a manner that makes it easy for others to observe what actions are being performed. When engaging customers, organizations may opt for transparency into operations (Buell et al. 2017, Buell, Porter, et al. 2020), costs (Mohan et al. 2020), sustainability initiatives (Buell and Kalkanci 2020, Kalkanci et al. 2018), governance practices (Kosack and Fung 2014) and more. Relatedly, information provision and information salience, both of which can be outcomes of transparency, are the acts of conveying new information (Cabral et al. 2010, Luca 2016), or increasing the prominence of disclosed information (Bollinger et al. 2011, Chetty et al. 2009, Luca and Smith 2013), respectively. In the bank-issued credit card context, where regulations stipulate that comprehensive information be made available about the terms and conditions of various offerings, tradeoff transparency may be most accurately characterized as increasing the salience of information about a product's tradeoffs, although numerous legal and online communication studies have demonstrated that very few customers read such statements (Ayres and Schwartz 2014, Obar and Oeldorf-Hirsch 2020). For example, in one study that tracked the behavior of 48,154 households across 90 websites, fewer than 0.2% accessed the terms of service, and most of those who access it read only a small portion (Bakos et al. 2014). Consequently, a service offering's downsides, promoted through tradeoff transparency, are more likely to be noticed, and in turn factored into a customer's purchase decision, than if those tradeoffs are solely presented in the text of the terms and conditions. In this section, we explore literature that informs how providing prospective customers with tradeoff transparency may affect customer acquisition and long-term engagement. Because of the conflicting findings among a portion of the relevant studies we cite, we have adopted the convention of stating non-directional hypotheses in null form and directional hypotheses in alternative form.

2.1 Tradeoff Transparency and Customer Acquisition

Research has highlighted how the revelation of unflattering information about a company's service offerings, such as poor quality or high prices, can reduce consumer demand. For example, research conducted in online platforms found that the first bad review posted about a firm's offerings can cost a company up to 13% of its revenue (Cabral et al. 2010) and that a one-star increase on Yelp can boost a company's revenue by 9% on average (Luca 2016). Chetty et al. (2009) found that making taxes salient in a retail context – that is, posting sales taxes along with the prices of goods – increases consumers' perceptions of the prices charged by the retailer, and reduces demand by 8%. Similarly, Finkelstein (2009) found that drivers who pay tolls electronically are substantially less aware of, and less sensitive to, toll rate increases, and Blake et al. (2017) illustrated that shrouding service fees on an online ticketing platform can substantially raise revenues. To the extent that disclosing sensitive information can adversely affect demand, these results suggest that firms intent on acquiring new customers may have powerful incentives to hide negative information from prospective buyers.

On the other hand, an extensive array of research in psychology and marketing has demonstrated the benefits of self-disclosure in fostering intimacy. Interpersonally, self-disclosure has been shown to be critical in developing social relationships (Jourard 1971) and in forging a heightened level of trust (Wheeless and Grotz 1977). People disclose in intimate relationships (Derlega 1984), as self-disclosure allows parties to gain knowledge about one another while reducing ambiguity about the other's intentions (Laurenceau et al. 1998, Perlman and Fehr 1987). Similar effects have been documented in non-interpersonal contexts. For example, research has shown how people become more attracted to, and more willing to engage with, computers that self-disclose sensitive information. If a computer shares that its "abilities are really limited…" in that it can "do word processing and spreadsheets, but it cannot do any kind of physical activity, like play sports or walk down the street," people report a higher level of attraction to it, and are more likely to engage with it (Moon 2000). Similarly, when companies self-disclose their costs of producing goods and services – a practice which implicitly reveals their profit margins – consumers report trusting the firm more, are more willing to engage with it, and sales increase (Mohan et al. 2020).

Furthermore, research on operational transparency has demonstrated how disclosing facets of the operation that are typically hidden from view – such as the work being performed behind the scenes – can enhance consumers' appreciation for the firm and their perceptions of the value it creates (Buell et al. 2017, Buell and Norton 2011). Moreover, such transparency can bolster trust in and engagement with the firm, even when it reveals imperfect performance (Buell, Porter, et al. 2020, Kalkanci et al. 2018). Common across all of these examples, however, is the notion of voluntary self-disclosure. Indeed, trust and engagement are not engendered when the revelation is involuntary (Hoffman-Graff 1977, Mohan et al. 2020). Taken together, the net effects on acquisition rates of a firm providing transparency into the tradeoffs inherent in its offerings are equivocal. The revelation of unflattering information may suppress demand, but voluntary self-disclosure may engender higher levels of trust and brand attraction. Due to these offsetting effects, we formulate the following hypothesis in null form:

Hypothesis 1 (H1): Providing prospective customers with tradeoff transparency will have no effect on the firm's rate of customer acquisition.

2.2 Tradeoff Transparency and Customer Engagement

It is well established that customer satisfaction plays a key role in a firm's long-term performance. Research has demonstrated that highly-satisfied customers are more loyal and more profitable over time (Anderson 1994, Anderson et al. 1994, Heskett et al. 1997, LaBarbera and Mazursky 1983). Satisfaction levels, in turn, are largely persistent between firms and individual customers, and are influenced in part by customer compatibility – the degree of fit between the needs of the customer and the capabilities of the operation (Buell, Campbell, et al. 2020). Customer compatibility can be influenced through a firm's segmentation strategies, wherein offerings are designed around the shared needs of particular operating segments of customers (Guajardo and Cohen 2017). However, the efficacy of such strategies is likely influenced by the degree of transparency firms provide to prospective customers into how those offerings are designed. Customers that are misaligned with a segmented service offering will be less satisfied with their experiences (Buell et al. 2016, Buell, Campbell, et al. 2020), and will impose more variability on the operation (Frei 2006) which, in turn, will hinder the firm's service performance (Karmarkar and Pitbladdo 1995, Chase and Tansik 1983).

Providing prospective customers with tradeoff transparency could improve customer compatibility in two ways: 1) by reducing Type I errors (false positives); and, 2) by reducing Type II errors (false negatives). In the context of customers choosing among service offerings, a Type I error occurs when a customer chooses an offering that is poorly aligned with his or her needs. When firms communicate the positive attributes of their offerings and omit the negative attributes, research in psychology has demonstrated that consumers often fill in the gaps in biased and overly favorable ways, and are more likely to choose the offering (Kivetz and Simonson 2000). Since such customers enter a service relationship with expectations that surpass the capabilities of the operation, the gap between their service expectations and experiences will result in dissatisfaction over time (McDougall and Levesque 2000, Oliver 1993, Parasuraman et al. 1985, Tse and Wilton 1988). In contexts with high switching costs, dissatisfied customers may remain with the firm (E. Anderson 1994, Buell et al. 2010, Yang and Peterson 2004), but are likely to spend less money than more highly-satisfied customers and may be more difficult and costly to serve (Coyles and Gokey 2005, Jones and Sasser 1995, Xue and Harker 2002). To the extent that providing prospective customers with tradeoff transparency can help reduce Type I errors, it stands to reason that the customers who choose the offering may be more aligned and satisfied with it, and may be more loyal and profitable to the firm.

On the other hand, in the context of customers choosing among service offerings, a Type II error occurs when a customer fails to select an offering that is well-aligned with his or her needs. Services have been characterized as "experience goods," since empirical research has demonstrated how customers are unable to fully assess them until after they have been delivered. Although providing consumers with incomplete information can trigger a psychological bias causing them to assume the best (Kivetz and Simonson 2000), when the lack of transparency is assumed to be volitional, people may assume the worst, undermining their trust and engagement (John et al. 2016). In this way, tradeoff transparency may also result in fewer Type II errors, further promoting service relationships that are value producing and satisfying for customers, and long-lived and highly profitable for firms (Heskett et al. 1997). To the extent that the information afforded by tradeoff transparency helps reduce Type I and Type II errors, we hypothesize that prospective customers who are provided with tradeoff transparency will select different offerings than those who are not:

> *Hypothesis 2 (H2):* Providing prospective customers with tradeoff transparency will result in the selection of different offerings than would be selected in the absence of such transparency.

We note that testing whether customers make *different* choices in the presence of tradeoff transparency, as hypothesized in H2, is not the same thing as testing whether customers make *better* choices in the presence of tradeoff transparency (i.e., whether customers select offerings that are better aligned with their needs and preferences). However, testing H2 is an important interim step, since if customer choices were not affected in any way by the presence of tradeoff transparency, then it could not follow that their choices were improved in its presence. Indeed, to the extent that the provision of tradeoff transparency reduces Type I and Type II errors, we would expect customer choices to not only be different, but to also be better in its presence than in its absence. Assessing the quality of individual customer choices on the basis of which offerings they choose could be compromised by the fact that customers may be making well-reasoned choices based on unobservable criteria. Indeed, part of what managers found attractive about the idea of tradeoff transparency, relative to other more prescriptive methods of influencing customer choices, is that it had the potential to empower customers, who have a much richer understanding of their own needs and preferences, to make better-informed decisions. In our experiment, we are able to directly observe how customers in the treatment and control conditions engage with the credit cards they choose over time, which serve as behavioral indicators of the long-run quality of their choices. We hypothesize that usage and retention will increase among customers who experience tradeoff transparency, since we predict that the offerings they will choose in the presence of tradeoff transparency will be better-aligned with their needs and preferences than the offerings they will choose in its absence:

Hypothesis 3 (H3): Providing prospective customers with tradeoff transparency will increase product usage among those who choose to select the offering, relative to not providing such transparency.

Hypothesis 4 (H4): Providing prospective customers with tradeoff transparency will increase product retention among those who choose to select the offering, relative to not providing such transparency.

Although we anticipate the main effects on usage and retention described above, prior research in the marketing domain has demonstrated that customers who are more familiar with a product category have the capacity to learn more from technical advertising (R. Anderson and Jolson 1980), and that more experienced consumers, by virtue of their familiarity, are better able to select information about attributes that are more predictive of product performance, in turn leading to better decisions that are more aligned with their needs (Johnson and Russo 1984). Hence, we predict that the long-run benefits of providing prospective customers with transparency into the tradeoffs of a service offering will accrue disproportionately to those with more prior category experience, who can learn from and make better use of the information transparency affords. If transparency can be best leveraged by more experienced customers to make choices that are better aligned with their needs, we would by extension predict that their levels of engagement would be disproportionately impacted. Accordingly, we hypothesize:

Hypothesis 5 (H5): The positive effect of tradeoff transparency on product usage will be most acute for customers who have more prior experience with the service category.

Hypothesis 6 (H6): The positive effect of tradeoff transparency on product retention will be most acute for customers who have more prior experience with the service category.

Lastly, a rich theoretical literature in economics and operations has modelled how customers sort among competing offerings, trading off price and service quality (Cachon and Harker 2002, Cohen and Whang 1997, Gabszewicz and Thisse 1979, Gans 2002, Hall and Porteus 2000, Karmarkar and Pitbladdo 1997, Li and Lee 1994, Shaked and Sutton 1982, Sutton 1986, Tirole 1990, Tsay and Agarwal 2000). Subsequent empirical work has documented that when service offerings are differentiated against the competitive set on the basis of quality, they attract more quality-sensitive customers, who are more likely to defect to higher-quality offerings. When offerings are instead differentiated on the basis of price, firms attract price-sensitive customers, who are less sensitive to quality, and more likely to defect to lower-priced offerings (Buell et al. 2016). Consistently, we predict that the effect of tradeoff transparency on customer engagement and retention will be negatively influenced by the presence of a promotion. Customers attracted to an offering on the basis of a discount or financial incentive may be less sensitive to the quality of the service experience than they are to the incentive itself, and hence, the benefits of increased alignment between their needs and the capabilities of the offering they select afforded by transparency are less likely to accrue to them. Moreover, to the extent that conflicts arise between the offering that's the best fit for a customer, and the promotional incentives provided to select it, we would further expect the presence of a promotion to diminish the impact of transparency on engagement and retention. Thus, we hypothesize:

Hypothesis 7 (H7): The positive effect of tradeoff transparency on product usage will be diminished among customers attracted to the offering by a promotion.

Hypothesis 8 (H8): The positive effect of tradeoff transparency on product retention will be diminished among customers attracted to the offering by a promotion.

3. Presentation of field experiment

To conduct this research, we partnered with Commonwealth Bank of Australia (CBA), which at the time of the experiment, had more than 1,000 branches, more than 45,000 employees, and more than 10 million retail banking customers. CBA was the largest issuer of credit cards in its market, with over 3 million credit card holders and annual transaction volume exceeding $50 billion USD. It offered nine types of personal credit cards, spread across three different families – awards cards, low rate cards, and low fee cards (**Figure 1.1**). We focus our analysis on customers who shopped for personal credit cards.

3.1 Data and methods

3.1.1 Participants, design, and procedure. From September 6, 2017 through February 4, 2018, we collaborated with CBA to conduct a field experiment, engaging all customers who visited the bank's credit card marketing website. 466,322 customers were randomly assigned to one of two experimental conditions.

Customers randomly assigned to the control condition, $TREAT_i=0$, observed a version of the website that was consistent with the bank's traditional marketing efforts – emphasizing the features and benefits of each credit card in its primary copy. Customers randomly assigned to the treatment condition, $TREAT_i=1$, observed an augmented version of the website, in which the primary copy additionally revealed the tradeoffs inherent in each offering (**Figure 1.2**). For example, customers in the control condition who shopped the low fee credit card would be able to view and learn about the card's benefits: $0 annual fee in the first year, free additional card holders, and up to 55 days interest-free on purchases. Furthermore, they would be able to view information about rates and fees, and about the security features of the card.

Customers randomly assigned to the treatment condition would additionally be able to view the tradeoffs inherent in the offering: (*i*) the purchase interest rate of 19.74% p.a. is higher than that of the low rate card, (*ii*) the low fee credit card does not include travel insurance, and (*iii*) there is an annual fee of $29 if the cardholder does not spend at least $1,000 on the card. The tradeoffs that were identified for display on the website were selected based on reviewing how the credit cards compared with one another, and finding the disadvantages of each, in terms of fees, limits, interest rates, and other features that had been identified by customers as important determinants of their service experiences interacting with the cards. We worked with the bank's copywriters, product, and compliance teams to arrive at the final copy that was used on the website, because it was important that the tradeoff transparency language complied with financial regulations, was devoid of jargon, was simple, and was brief.

Figure 1.1. Credit card offerings from CBA during the time of the experiment. The bank offered cards in three families: 1) low rate cards, 2) low fee cards, and 3) awards cards. Because the credit card offerings were designed to serve customers with varying needs and preferences, the value proposition of each credit card included both benefits and tradeoffs, making it more appropriate for some prospective customers and less appropriate for others.

	Benefits (included in treatment and control)	**Tradeoffs** (included in the treatment only)
Low Rate	• Our lowest interest rate on purchases, currently 13.24% p.a. • Minimum credit limit of $500 • Free additional card holder	• The annual fee of $59 is higher than our Low Fee credit card • Does not earn awards points • Does not include travel insurance • International purchases may incur international transaction fees
Low Rate Gold	• Our lowest interest rate on purchases, currently 13.24% p.a. • International travel insurance included • Minimum credit limit of $4,000 • Free additional card holder	• There's an annual fee of $89 • Does not earn awards points • International purchases may incur international transaction fees
Low Fee	• $0 annual fee in the first year* and following years when you spend at least $1,000 in the previous year • Free additional card holder	• Does not earn awards points • Does not include travel insurance • There's an annual fee of $29 in the second and later years if you spend less than $1,000* in the previous year • International purchases may incur international transaction fees
Low Fee Gold	• $0 annual fee in the first year* and following years when you spend at least $10,000 in the previous year • International travel insurance included • Free additional card holder	• The purchase interest rate of 19.74% p.a. is higher than Low Rate credit cards • Does not earn awards points • There's an annual fee of $89 in the second and later years if you spend less than $10,000* in the previous year • International purchases may incur international transaction fees
Student	• $0 annual fee in the first year** • Free additional card holder	• The purchase interest rate of 19.74% p.a. is higher than Low Rate credit cards • Does not earn awards points • Does not include travel insurance • International purchases may incur international transaction fees
Awards	• Earn up to 1.5 Awards points for every dollar you spend • Access to the largest rewards program of any bank in Australia*	• The purchase interest rate of 20.24% p.a. is higher than Low Rate and Low Fee credit cards • There's an annual fee of $59 • Does not include travel insurance • International purchases may incur international transaction fees • There's a $10 p.a. additional cardholder fee
Gold Awards	• Earn up to 2 Awards points for every dollar you spend • Access to the largest rewards program of any bank in Australia* • International travel insurance included when you activate cover within [our online banking platform]	• The purchase interest rate of 20.24% p.a. is higher than Low Rate and Low Fee credit cards • There's an annual fee of $119 • International purchases may incur international transaction fees • There's a $10 p.a. additional cardholder fee
Platinum Awards	• Earn up to 2.5 Awards points for every dollar you spend • Access to the largest rewards program of any bank in Australia* • International travel insurance included when you activate cover within [our online banking platform] • No international transaction fees on overseas purchases made in-store and online using your CommBank Platinum American Express card	• The purchase interest rate of 20.24% p.a. is higher than Low Rate and Low Fee credit cards • There's an annual fee of $249 • There's a $10 p.a. additional cardholder fee
Diamond Awards	• Earn up to 3 Awards points for every dollar you spend • Access to the largest rewards program of any bank in Australia* • International travel insurance included when you activate cover within [our online banking platform] • No international transaction fees on overseas purchases made in-store and online using your CommBank Diamond American Express card	• The purchase interest rate of 20.24% p.a. is higher than Low Rate and Low Fee credit cards • There's an annual fee of $349 • There's a $10 p.a. additional cardholder fee

Figure 1.2. Website copy used in the treatment and control conditions for each card. Customers randomly assigned to the control condition only experienced the marketing of each credit card's benefits, which is common practice in the industry. Customers randomly assigned to the treatment condition were additionally presented with transparency into the tradeoffs of each offering.

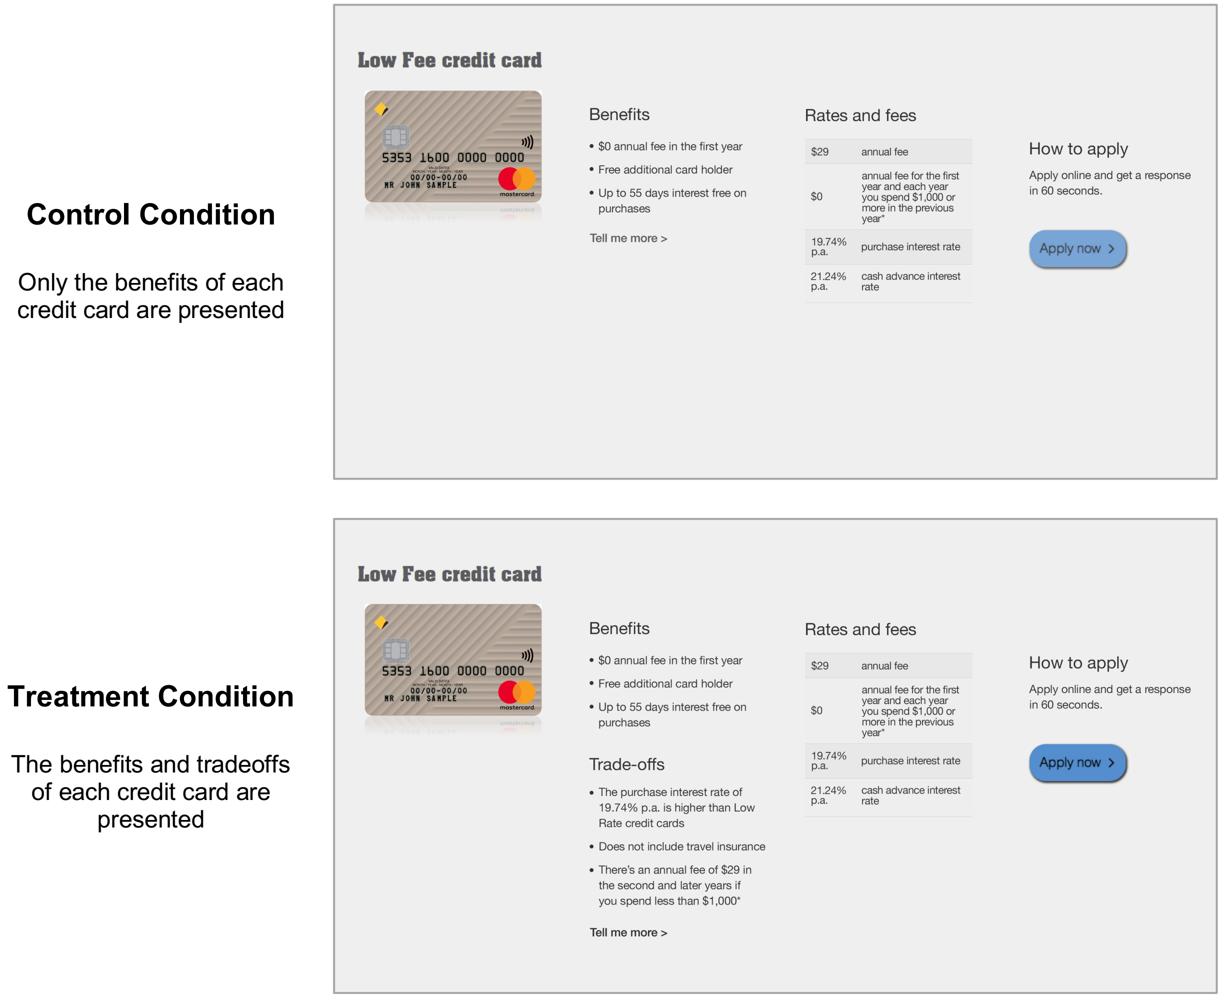

Figure 1.3. Example credit card product detail pages in the control and treatment conditions. The experimental manipulation involved more than 50 blocks of content on more than 20 pages of CBA's public-facing and secure online banking websites, such that every credit card marketing webpage where the features and benefits of a credit card were described, so too were its tradeoffs for customers randomly assigned to the treatment.

We note that in some cases, tradeoffs were presented in absolute terms, such as "there's a $10 p.a. additional cardholder fee," whereas in other cases, when a feature compared unfavorably with one or more cards, the copy was written to reveal the relative comparison, such as "the annual fee of $59 is higher than our Low Fee credit card." Importantly, as in many other developed economies, the Australian financial services industry is highly regulated. Banks are required by law to disclose the full terms and conditions of each offering in product disclosure statements, which are continuously monitored and assessed by the Australian Securities and Investments Commission (ASIC), the country's integrated corporate, markets, financial services, and consumer credit regulator. These statements, which appeared toward the bottom of each page of CBA's credit card website, disclosed information about the tradeoffs of each offering to customers in both conditions. However, the experimental treatment promoted its salience by moving its messaging into the page's primary copy. In this way, the experiment was designed to cleanly identify the effects of increasing the transparency of information about an offering's tradeoffs, on the acquisition rates, choices, and subsequent engagement, of prospective customers (**Figure 1.3**).

During the time of our experiment, the bank offered three channels through which customers could access information about its credit cards: a public-facing website, a secure online banking website, and a mobile banking application. Our experiment manipulated the presentation of information on the two websites, which historically had accounted for roughly 80% of new credit card applications that came in through digital channels. The mobile application, in which it was infeasible to experimentally manipulate content, contained an abbreviated version of the information displayed online to the customers in the control condition. This content was held constant during the period of our experiment, and as such, it should not serve as a confound in our analyses.

For each customer, assignment to an experimental condition was randomly determined at the beginning of their first visit to the credit card section of one of the bank's two websites. The consistency of a customer's random assignment to an experimental condition across channels was controlled by a cookie that was passed by the bank's public website to each customer's browser. When customers visited the online banking website, their assigned conditions were saved in the bank's databases. Although this design largely ensured consistent presentation of content across multiple sessions, 44,765 customers were identified to have experienced multiple conditions and were dropped from the analysis. We also dropped an additional 20,008 customer observations that were missing relevant data elements – such as the date the customer visited the credit card website, or the types of credit cards for which they browsed – that prevented tracking the customer's progression through the acquisition funnel. We additionally dropped 1,514 customers from the analysis who ultimately applied outside the experimental period. Moreover, midway through the experiment, the bank detected inconsistent benefits information across the treatment and control conditions for one type of low fee card on one page of the bank's website. To distinctly identify the effect of the treatment on the treated customers, for our primary analysis, we withhold observations from the 6,999 customers in both conditions who visited this specific page, though we note in an analysis in the online appendix that the results are substantively similar with these customers included. Our primary analysis, therefore, includes observations from a total of 393,036 customers, 195,438 of whom were randomly assigned to the control condition, and 197,598 of whom were randomly assigned to the treatment condition. **Table 1.1** presents summary statistics for customers in both experimental conditions on a variety of observable dimensions. We control for these customer-level factors in our empirical models, though we note that customers were well-balanced on both dimensions prior to random assignment, and as we show in analyses presented in the online appendix, results are substantively similar if controls are not included. These control variables were captured as of each customer's first documented arrival on the credit card website, and thus, were unaffected by the experimental treatment. Customer-level data captured temporally after the random assignment and presentation of the experimental treatment, such as which credit card each customer selected, are themselves outcomes in the notional experiment, and thus, would be considered "bad controls" if included in our analysis (Angrist and Pischke 2008). Accordingly, all treatment effects presented in this paper can be interpreted as causal effects.

	All			Control			Treatment			Diff.
	N	Mean	SD	N	Mean	SD	N	Mean	SD	P-value
Arrived during the promo period	393,036	63.00%	0.48	195,438	62.89%	0.48	197,598	63.11%	0.48	0.16
Customer demographics										
Customer age	392,942	40.12387	15.59	195,386	40.12955	15.62	197,556	40.11824	15.57	0.82
Customer tenure (months)	392,357	198.5407	138.62	195,092	198.3877	138.74	197,265	198.692	138.50	0.49
Male indicator	392,863	45.03%	0.50	195,350	45.11%	0.50	197,513	44.95%	0.50	0.29
Product indicators										
Transaction product	392,874	89.73%	0.30	195,355	89.72%	0.30	197,519	89.73%	0.30	0.91
Savings product	392,874	67.72%	0.47	195,355	67.71%	0.47	197,519	67.74%	0.47	0.89
Home loan product	392,874	22.03%	0.41	195,355	22.06%	0.41	197,519	22.01%	0.41	0.72
Home insurance policy	392,874	11.49%	0.32	195,355	11.51%	0.32	197,519	11.47%	0.32	0.68
Personal loan product	392,874	10.01%	0.30	195,355	9.96%	0.30	197,519	10.06%	0.30	0.30
Retirement product	392,874	4.75%	0.21	195,355	4.77%	0.21	197,519	4.73%	0.21	0.53
Motor insurance policy	392,874	4.04%	0.20	195,355	4.01%	0.20	197,519	4.07%	0.20	0.34
Term deposit product	392,874	3.89%	0.19	195,355	3.90%	0.19	197,519	3.88%	0.19	0.71
Credit card product	392,874	62.84%	0.48	195,355	62.90%	0.48	197,519	62.78%	0.48	0.45

Table 1.1. Summary statistics for customer characteristics in the control and treatment conditions. Customer characteristics were captured as of the time of each customer's first documented arrival on the credit card website, and thus, are unaffected by the experimental treatment. We control for all of these customer-level factors in our empirical models, though we note that customers in both conditions were quite well-balanced on these observable dimensions prior to their random assignment, and that uncontrolled analyses, presented in the online appendix, reveal substantively similar results. Differing numbers of observations for different variables reflect missing data.

Furthermore, on November 1, 2017 (midway through our period of experimentation), the bank publicly announced a promotion, offering \$300 cash back to customers who opened a credit card in the Low Rate family and spent \$1,000 in purchases with the new card within the first 90 days of activation. The announcement was broadly supported by advertising (e.g., television, radio, and print). Random assignment to the experimental conditions proceeded during this promotional period, affording an opportunity to assess whether providing customers with incentives to choose particular offerings moderates the effect of presenting tradeoffs on customer acquisition and engagement. The promotional period ran from November 1 through the end of the experiment. We identified 145,412 customers who initially visited the credit card website prior to the November 1 promotion announcement to not be treated with the promotion, $PROMO_i$ = 0, and 247,624 customers who initially visited on November 1 or after to be treated with the promotion, $PROMO_i$ =1. We cannot directly observe whether customers who initially visited during this period were actually exposed to the promotion, so results from this variable should be interpreted according to the intention-to-treat principle. Our analysis yields an unbiased estimate of the effect of the promotion on the behavior of prospective customers, at the average level of exposure to the promotion.

3.1.2 Acquisition measures. In order to assess the impact of providing transparency into the tradeoffs of a service offering on rates of customer acquisition we tracked: (1) whether the customer progressed through each stage of the acquisition funnel, $\Pr\left(ACQ_i\right)$, and (2) which types of credit card accounts customers opened, $\Pr\left(CHOICE_i\right)$. At the time of our study, the credit card acquisition funnel for our partner bank had six stages. In order to gain a detailed understanding of how the experimental treatment affected customer acquisition, we used six binary variables to capture each customer's progress through the acquisition funnel. In particular, we coded each stage of the acquisition funnel to be 1 if the customer passed the stage, and 0 if she did not.

The first stage of the acquisition funnel occurred when the customer *started an application*. During the application process, customers of the partner bank answered a series of 10 to 40 questions, providing their demographic information as well as their financial information (such as assets, liabilities, income, and expenses). For the more than 95% of prospective customers who had existing accounts with the bank, this stage of the process was streamlined considerably; by simply logging into their online banking accounts, these customers would be able to transfer their information from their existing accounts to their credit card applications. Answering these questions typically took customers 5-15 minutes, depending on how much information was required to complete their financial profile.

Submitting the answers to these questions triggered the second stage of the acquisition funnel, which was known as a *soft submission*. In this phase, the bank parsed through the information provided using an automated process to conduct a risk assessment and determine whether the customer's financial profile warranted approval for the credit card for which she had applied, as well as her eligible credit limit. After the results of this process were shared with the customer, the customer could choose to move on to the third stage, providing a *hard submission.* The application submitted at this stage constituted a formal application submission for a credit card; at this stage, customers provided verification documents required by the bank to substantiate their income as well as financial holdings and obligations that resided outside the bank. This stage could last two weeks or longer, depending on the amount of information provided by the customer and the amount of time the customer took to provide all the necessary information.

If the bank's process during the hard submission stage substantiated the information provided in the application, the customer was then moved into the fourth stage, wherein her account was *opened.* Here, the bank created and issued a new credit card for the customer and mailed it to her preferred address. This process typically took about a week. To test whether providing transparency influenced customer choices, we take a snapshot at this stage, comparing which types of card accounts were opened by customers randomly assigned to the treatment and control conditions.

In the fifth stage, the customer *activated* the card either by phone, online, through the mobile application, or by inserting the card into one of the bank's ATM machines. Finally, in the sixth stage, the customer *used* the card to make her first purchase; customers that reached the final stage of the acquisition funnel comprised 4.05% of all customers who shopped for a credit card. For our analyses, we denote a customer to have reached this final stage if a purchase was made with the card within two billing cycles of its activation. Imposing this constraint facilitates comparability among customers acquired early and late in our experimental period, as failure to do so would afford those who opened their accounts early more time to demonstrate usage. We note that 95.42% of all cards that were eventually used to make a purchase in our dataset were used within two billing cycles of their initial activation, and that all results presented in this manuscript are substantively similar if usage is instead defined by whether any transaction behavior is ever observed in our data.

Customer progress through the acquisition funnel could stall during any of these six stages, as each stage was comprised of different dynamics with respect to the customer's experience and the bank's operating costs. Hence, analyzing each stage separately provides important insights about the costs and benefits of providing prospective customers with transparency into the tradeoffs inherent in a service offering.

3.1.3 Engagement measures. For customers who made it all the way through the acquisition funnel, we tracked customer engagement with the card in two ways. First, we tracked product usage by examining logged monthly spending during the first nine months the customer held the card, $\ln\left(SPEND_{it}\right)$. Spend is a relevant engagement metric, both because it is a direct measure of usage that is an important behavioral indicator of a card's utility for the customer, and because it is an important driver of profitability for the bank. Bank-issued credit cards generate revenue through interest spreads (interest paid by customers who don't fully pay off their balances at the end of the month), interchange fees paid by merchants (small percentages of each transaction), and customer fees (usage fees, service fees, and penalty fees). Customer spend directly influences the bank's first two revenue streams.

Second, we tracked customer retention, i.e., whether the customer closed the credit card account during the first nine months it was open, $\Pr\left(RETAIN_{i}\right)$. As with spend, retention is a relevant measure of customer engagement, because loyalty is an important behavioral indicator of the card's utility to the customer, and because it is an important driver of profitability (Levesque and McDougall 1996). Previous research conducted in banking demonstrated that reducing defections by just 5% can increase firm profitability by as much as 100% (Reichheld and Sasser 1990).

We note that these engagement measures were queried on August 28, 2018, which means that six full months of account data are available for any customer who activated their credit cards before February 28, 2018, and nine full months of account data are available for any customer who activated their credit cards before November 30, 2017. Importantly, all customers who were denoted to have used their cards in our analysis had at least six months of available account data. Consistently, we track spend during any observed month for which a credit card account was open and activated, and we track retention six and nine months into the customer relationship.

3.1.4 Control measures. Although customers randomly assigned to the treatment and control conditions appear quite well balanced on observable dimensions (**Table 1.1**), to get the cleanest possible estimates of the effects of the experimental treatment on treated customers, we control for a vector of customer-level factors which were observable prior to each customer's random assignment. These factors include the non-linear effect of the customer's age, AGE_i and AGE_i^2, and tenure with the bank (in months), $TENURE_i$, and $TENURE_i^2$, their gender $GENDER_i$ (operationalized as an indicator for whether the customer was male), and indicators for whether they held a home loan product, $HLOAN_i$, a personal loan product, $PLOAN_i$, a transaction product, $TRANS_i$, a savings product, SAV_i, a home insurance policy, $HINS_i$, a vehicle insurance policy, $VINS_i$, or a retirement product with the bank, RET_i. For the acquisition analyses, we also include a vector of date indicator variables, X_i, signifying the first week the customer visited the credit card pages on the bank's website, to control for time-varying factors that may have affected each customer's initial motivation for seeking a credit card. For the engagement analyses, we include a similar vector of date indicator variables, X_i, signifying the first week the customer activated their credit card, to control for time-varying factors that may influence the way customers interact with the credit card (e.g., a credit card activated before the holidays may be used more intensively than one activated after, etc.)

3.1.5 Category experience and financial incentives. Lastly, to test Hypotheses 5-8, that the effects of providing transparency into the tradeoffs inherent in a service offering will accrue disproportionately to customers with more experience with the product category (H5-H6), and will have a greater effect on customers who were not attracted by a financial incentive (H7-H8), we incorporate data on proxies for both factors into our analysis. First, since we lack complete data on the degree of credit card experience for each prospective customer (e.g., because not every customer provides their financial history, and many customers may have accounts with financial institutions that are not our partner bank), we use the prospective customer's age as the primary proxy in our analysis. In the market where CBA operated during the time of our experiment, the minimum age to apply for a credit card was 18. As we demonstrate graphically in the online appendix, data captured by CBA reveals that the probability a prospective customer has at least one credit card increases concavely in age. As such, data on the ages of credit card browsers and applicants during our study was right skewed. The average credit card browser was 40.12 years old and the median browser was 37 years old. The average credit card applicant was 31.27 years old, and the median applicant was 28 years old. To test hypotheses 5 and 6, that the effects of providing prospective customers with transparency into the tradeoffs of each offering will have a disproportionate effect on the product usage and retention of customers with more prior category experience, we complement our analyses on the full sample of data with split sample analyses of customers who were 28 and younger (less experienced with credit cards), and customers who were older than 28 (more experienced with credit cards). Using the median applicant age to conduct a split sample analysis yields a comparable number of more-experienced and less-experienced customers. However, as we document in the online appendix, our results are substantively similar with different cutoffs, for example, splitting the bottom quartile of the credit card applicant data (24 and younger) from the top three quartiles.

Second, to test whether our experimental treatment had differential effects for customers who were motivated to shop for a credit card by a promotion, we include an interaction term in all of our empirical specifications, $PROMO_i$, denoting whether the customer first visited the credit card website or applied for a credit card on or after November 1, 2017, the day the promotion was publicly announced. This empirical approach enables us to directly test the significance of the effect of the treatment on customers who conducted their credit card searches organically (e.g., before November 1, in the absence of a promotion), and also to test whether the treatment had a differential effect among customers who are more likely to have been influenced by the promotion (e.g., on or after November 1, initiating their search or applying for a credit card while the promotion was live). Since we cannot observe whether customers who first began their search on or after November 1, 2017 were directly influenced by the promotion, this empirical approach constitutes an intention to treat design, and hence, should be interpreted as a conservative estimate of the contingent effect of the promotion on the transparency treatment. We note that this empirical design is consistent with prior research in the marketing literature that has analyzed customers who were acquired by means of a promotion differently from those who were acquired organically (Anderson and Simester 2004, Krishnamurthi and Raj 2006, Lewis 2006).

3.2 Empirical approach

We analyze rates of acquisition, card choice, spend, and customer retention as a part of our analysis. This section outlines our empirical approach for each analysis.

3.2.1 Acquisition funnel. To estimate the effects of the experimental treatment on the probability that customer *i* reached each stage of the acquisition funnel (application started, soft submission completed, submission completed, account opened, card activated, card used), ACQ_i, we use the following logistic regression model, with standard errors clustered by the date she first visited the credit card marketing site. Clustering standard errors in this way is intended to account for the possibility that customers who first arrive on the same day may be related to each other, for example in the driver of their motivation to consider a credit card (e.g., were attracted to the site by the same advertisement, were motivated to seek a credit card for the same upcoming holiday or event, etc.). We note that ACQ_i for any particular customer refers to the furthest acquisition funnel stage she reached for any credit card for which she applied during the experimental period.

$$\Pr\left(ACQ_i\right) = f\left(\begin{gathered} \alpha_0 + \alpha_1 TREAT_i + \alpha_2 PROMO_i + \alpha_3 TREAT_i \times PROMO_i + \\ \alpha_4 AGE_i + \alpha_5 AGE_i^2 + \alpha_6 TENURE_i + \alpha_7 TENURE_i^2 + \alpha_8 GENDER_i + \\ \alpha_9 HLOAN_i + \alpha_{10} PLOAN_i + \alpha_{11} TRANS_i + \alpha_{12} SAV_i + \alpha_{13} HINS_i + \\ \alpha_{14} VINS_i + \alpha_{15} RET_i + X_i + \epsilon_i \end{gathered}\right) \quad (1)$$

3.2.2 Card choice. To estimate the effects of the experimental treatment on the probability that customer *i* opened each type of card (awards, diamond, gold low fee low fee gold, low rate, low rate gold, platinum, or student), $CHOICE_i$, we use the following logistic regression specification with standard errors clustered by the date she first visited the credit card marketing site:

$$\Pr\left(CHOICE_i\right) = f\left(\begin{array}{c} \beta_0 + \beta_1 TREAT_i + \beta_2 PROMO_i + \beta_3 TREAT_i \times PROMO_i + \\ \beta_4 AGE_i + \beta_5 AGE_i^{\,2} + \beta_6 TENURE_i + \beta_7 TENURE_i^2 + \beta_8 GENDER_i + \\ \beta_9 HLOAN_i + \beta_{10} PLOAN_i + \beta_{11} TRANS_i + \beta_{12} SAV_i + \beta_{13} HINS_i + \\ \beta_{14} VINS_i + \beta_{15} RET_i + X_i + \epsilon_i \end{array}\right) \quad (2)$$

To account for time-varying differences in the motivation to pursue a credit card, which could affect both persistence in the acquisition process, and one's choice of credit card, we include a vector of indicator variables, X_i, in Models (1) and (2) denoting the week the customer first visited the marketing website.

3.2.3 Spend. To estimate the effects of the experimental treatment on the spend of customer *i*, in month *t,* observed in months after the activation of her credit card, we use the following fixed effects panel model, with standard errors clustered by the date she activated the card. Clustering standard errors in this way is intended to account for the possibility that customers who activate their cards on the same day may exhibit similar spending patterns (e.g., customers who activate immediately before a holiday may spend more than those who activate immediately after, etc.).

$$\ln\left(SPEND_{it}\right) = f\left(\begin{array}{c} \gamma_0 + \gamma_1 TREAT_i + \gamma_2 PROMO_i + \gamma_3 TREAT_i \times PROMO_i + \\ \gamma_4 AGE_i + \gamma_5 AGE_i^{\,2} + \gamma_6 TENURE_i + \gamma_7 TENURE_i^2 + \gamma_8 GENDER_i + \\ \gamma_9 HLOAN_i + \gamma_{10} PLOAN_i + \gamma_{11} TRANS_i + \gamma_{12} SAV_i + \gamma_{13} HINS_i + \\ \gamma_{14} VINS_i + \gamma_{15} RET_i + Z_i + \epsilon_{it} \end{array}\right) \quad (3)$$

3.2.4 Retention. To estimate the effects of the experimental treatment on the retention of customer *i,* six and nine months after the activation of her card, we use the following logistic regression, with standard errors clustered by the date she activated the card.

$$\Pr\left(RETAIN_i\right) = f\left(\begin{array}{c} \delta_0 + \delta_1 TREAT_i + \delta_2 PROMO_i + \delta_3 TREAT_i \times PROMO_i + \\ \delta_4 AGE_i + \delta_5 AGE_i^2 + \delta_6 TENURE_i + \delta_7 TENURE_i^2 + \delta_8 GENDER_i + \\ \delta_9 HLOAN_i + \delta_{10} PLOAN_i + \delta_{11} TRANS_i + \delta_{12} SAV_i + \delta_{13} HINS_i + \\ \delta_{14} VINS_i + \delta_{15} RET_i + Z_i + \epsilon_t \end{array}\right) \quad (4)$$

To account for time-varying differences in credit card usage in Models (3) and (4), we include a vector of indicator variables, Z_i, denoting the calendar week in which the customer activated the card.

3.3 Analysis and results

3.2.1 Acquisition funnel (H1). **Figure 1.4** graphically displays the marginal effects of the treatment on the acquisition funnel of the treated customers. The results show three features of consequence. First, rates of conversion were higher during the promotion period, likely owing to the efficacy of the promotion in attracting interested customers to the website and in enhancing the appeal of the promoted products. Second, during the promotion period, when offered transparency into the tradeoffs inherent in each credit card, significantly fewer customers chose to start the application process (17.78% vs 17.33%, p<0.01) or soft submit their application (14.53% vs 14.17%, p<0.01). Interestingly, these differences did not emerge among customers who shopped for credit cards outside the promotion period. Finally, and most importantly, during both periods, there were no significant differences in rates of conversion through the various phases of the acquisition funnel during the submission phase and beyond. As we document in the online appendix, we note that these patterns are substantively similar among customers with high and low experience levels with credit cards. Consequently, we fail to reject H1. Taken together, these results suggest that from an acquisition perspective, voluntarily providing prospective customers with transparency into the tradeoffs of the firm's offerings had an insignificant effect on the overall rate of customer acquisition.

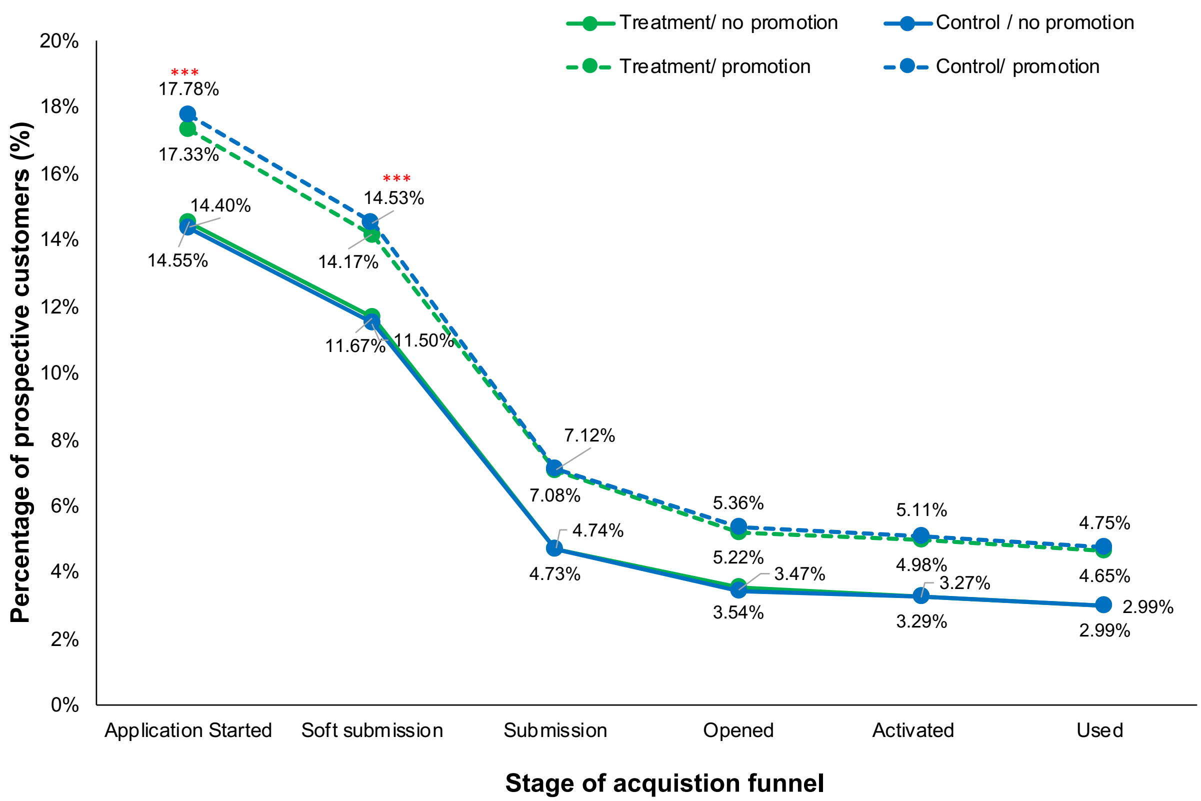

Figure 1.4: Acquisition funnel conversion percentages during the promotion and non-promotion periods. Marginal effects are from logistic regression models, estimated with robust standard errors clustered at the first website visit date level. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively. The results indicate that conversion rates are significantly higher during the promotion period. Moreover, although fewer customers who are exposed to the treatment during the promotion period start or soft submit an application, the treatment ultimately had an insignificant effect on rates of customer acquisition, neither during the promotion nor the non-promotion periods. Consequently, we fail to reject H1.

3.2.2 Card choice (H2). **Table 1.2**, Panel A documents the marginal effects of Model (2), regressed first on the full sample of customers who opened credit cards during the entire experimental period, and then on a split sample, median-split by age, with customers at and below the median age presented first and customers above the median age presented second. This analysis is replicated on the subsample of customers who arrived during the non-promotion period and promotion period in Panels B and C, respectively. Each table directly compares the probabilities that applicants in the control and treatment conditions would open different types of cards. The results demonstrate that, consistent with H2, the treatment did affect the credit card choices of customers who opened credit card accounts. Panel A demonstrates that, controlling for other factors, customers in the treatment condition were less likely to open the low rate card (β = -0.144, $p<0.05$) and more likely to open the low rate gold card (β = 0.212, $p<0.05$) and platinum card (β = 0.266, $p<0.05$) than customers in the control condition.

A. Entire period (n=18,152)

	All customers (n=18,152)				Customer age ≤ 28 (n=9,090)				Customer age > 28 (n=9,062)			
	Control	Treatment		P-value	Control	Treatment		P-value	Control	Treatment		P-value
Awards card	4.81%	4.97%		0.908	4.48%	5.16%		0.992	5.14%	4.76%		0.915
Diamond card	2.16%	2.26%		0.396	1.66%	2.12%		0.883	2.64%	2.46%		0.239
Gold card	2.75%	2.77%		0.195	2.83%	2.99%		0.151	2.71%	2.52%		0.664
Low fee card	14.78%	14.54%		0.324	16.08%	15.11%		0.715	13.43%	14.02%		0.234
Low fee gold card	2.76%	3.07%		0.977	2.85%	3.09%		0.974	2.68%	3.04%		0.936
Low rate card	51.75%	49.81%	**	0.033	51.00%	48.48%	*	0.093	52.57%	51.11%	*	0.098
Low rate gold card	11.96%	12.96%	**	0.013	10.60%	11.50%		0.111	13.32%	14.42%	*	0.092
Platinum card	5.41%	6.01%	**	0.013	3.98%	4.89%	**	0.045	6.85%	7.14%		0.278
Student card	3.56%	3.65%		0.938	6.73%	6.99%		0.818	0.78%	0.71%		0.418

B. Non-promotion period (n=5,087)

	All customers (n=5,087)				Customer age ≤ 28 (n=2,632)				Customer age > 28 (n=2,455)			
	Control	Treatment		P-value	Control	Treatment		P-value	Control	Treatment		P-value
Awards card	6.14%	6.09%		0.944	5.71%	5.62%		0.912	6.75%	6.72%		0.979
Diamond card	3.15%	2.81%		0.437	2.55%	2.44%		0.852	3.81%	3.25%		0.341
Gold card	3.85%	3.21%		0.209	4.03%	2.88%		0.148	3.88%	3.54%		0.641
Low fee card	18.56%	19.48%		0.342	18.48%	19.09%		0.652	18.54%	20.22%		0.219
Low fee gold card	4.19%	4.17%		0.977	3.90%	3.88%		0.974	4.51%	4.56%		0.952
Low rate card	43.29%	39.72%	**	0.027	43.88%	40.67%	*	0.094	42.71%	38.51%	*	0.068
Low rate gold card	10.54%	12.79%	**	0.010	9.15%	11.15%	*	0.099	12.17%	14.79%	*	0.062
Platinum card	5.80%	7.47%	**	0.011	4.37%	6.35%	**	0.040	7.52%	8.86%		0.265
Student card	5.41%	5.42%		0.989	9.18%	9.55%		0.745	1.31%	0.87%		0.365

C. Promotion period (n=13,065)

	All customers (n=13,065)				Customer age ≤ 28 (n=6,458)				Customer age > 28 (n=6,607)			
	Control	Treatment		P-value	Control	Treatment		P-value	Control	Treatment		P-value
Awards card	4.29%	4.56%		0.438	4.02%	5.00%	*	0.065	4.57%	4.11%		0.358
Diamond card	1.80%	2.06%		0.298	1.37%	2.04%	**	0.047	2.20%	2.19%		0.957
Gold card	2.34%	2.62%		0.292	2.41%	3.11%	*	0.056	2.36%	2.22%		0.702
Low fee card	13.33%	12.66%		0.291	15.12%	13.55%	*	0.093	11.56%	11.82%		0.738
Low fee gold card	2.24%	2.66%		0.108	2.47%	2.77%		0.456	2.04%	2.54%		0.199
Low rate card	55.09%	53.70%		0.130	53.91%	51.66%	*	0.094	56.31%	55.65%		0.598
Low rate gold card	12.55%	13.07%		0.377	11.23%	11.69%		0.573	13.83%	14.43%		0.537
Platinum card	5.27%	5.51%		0.578	3.91%	4.48%		0.239	6.65%	6.58%		0.917
Student card	3.43%	3.59%		0.618	5.87%	6.15%		0.639	1.02%	1.10%		0.834

Table 1.2: Effects of the treatment on card choice. Marginal effects from logistic regression models, estimated with robust standard errors clustered at the first website visit date level are presented. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively. The results indicate that providing transparency into the tradeoffs of the service offering influenced credit card choices, which is consistent with H2.

Examining the split sample analyses provides evidence that transparency into the tradeoffs of each credit card influenced the behavior of customers with varying levels of experience differently. For example, less experienced customers exposed to the treatment were more likely to open the platinum card ($\beta = 0.393$, $p<0.05$) and were marginally less likely to open the low rate card ($\beta = -0.133$, $p<0.10$). More experienced customers exposed to the treatment were marginally less likely to open the low rate card ($\beta = -0.162$, $p<0.10$), and were marginally more likely to open the low rate gold card ($\beta = 0.208$, $p<0.10$). These differences are interesting, since they suggest that different types of customers with different needs evaluated the tradeoffs differently, which in turn, differentially affected their choices. Moreover, these diverging effects of the treatment among customers with varying experience levels further suggest that the results were not merely driven by tradeoff copy that made some credit cards appear clearly dominant to others.

As additional evidence, it is revealing to consider the differential effects of the treatment on the choices of customers who applied for credit cards in the absence and presence of a promotion, which are presented in Panels B and C, respectively. For customers who arrived during the non-promotion period, the patterns of results among all, less experienced, and more experienced customers were consistent with those described above across the entire experimental period, with the addition that less experienced customers were marginally more likely to select the low rate gold card (β = 0.225, p<0.10) when provided transparency into each card's tradeoffs. By contrast, during the promotion period, which likely attracted customers with different needs and preferences, the treatment dissuaded less experienced customers from the low rate (β = -0.092, p<0.10) and low fee cards (β = -0.129, p<0.10), and increased their interest in the awards (β = 0.229, p<0.10), diamond (β = 0.422, p<0.05), and gold cards (β = 0.265, p<0.10). Interestingly, during the promotion period, the treatment exhibited no effect on the credit card choices of more experienced customers. In sum, these results suggest that, consistent with H2, transparency influenced customers' credit card choices.

As noted previously, an alternative question is whether tradeoff transparency influenced customers' choices for the better. In the online appendix, we provide some evidence based on observable customer characteristics that the choices customers made in the presence of tradeoff transparency may have been more compatible with their banking needs. For example, customers with fewer banking products, who may have been more likely to use their credit cards for short-term borrowing, may have been best matched with the low rate card, which offers the most favorable interest rate. We find that such customers were marginally more likely to choose the low rate card in the presence of tradeoff transparency (γ=-0.069, p<0.10); a tendency that was especially strong during the non-promotion period (γ=-0.179, p<0.01).

Interestingly, although younger customers are more likely to choose the student card in general (β =-0.225, p<0.01), we do not find that their propensity to do so is influenced by the presence of tradeoff transparency (β =0.008, p=*N.S.*). These results may suggest that when the intended segment for an offering is already clearly communicated (e.g., the student card is for students), the addition of tradeoff transparency may be less impactful. A complementary potential explanation is that, consistent with the arguments in Section 2.2, customers with less prior category experience may be less impacted by tradeoff transparency.

	(1) Ln(Spend)	(2) Ln(Spend)	(3) Ln(Spend)	(4) Ln(Spend)	(5) Ln(Spend)	(6) Ln(Spend)
Treatment	0.080	0.035	0.138**	0.095*	0.033	0.176***
	(0.052)	(0.085)	(0.069)	(0.054)	(0.092)	(0.062)
Promotion	-0.104	0.019	-0.230	-0.371***	-0.150	-0.593***
	(0.103)	(0.131)	(0.151)	(0.114)	(0.125)	(0.166)
Treatment x promotion	-0.097	-0.054	-0.162**	-0.109*	-0.045	-0.199**
	(0.061)	(0.103)	(0.082)	(0.064)	(0.110)	(0.080)
Customer age	0.072***	0.077	0.009	0.046***	0.137	0.027
	(0.008)	(0.129)	(0.018)	(0.009)	(0.136)	(0.019)
Customer age²	-0.001***	-0.001	-0.000	-0.001***	-0.002	-0.000
	(0.000)	(0.003)	(0.000)	(0.000)	(0.003)	(0.000)
Customer tenure	0.002***	0.005***	-0.001	0.000	0.004***	-0.002***
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Customer tenure²	-0.000	-0.000***	0.000*	0.000	-0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.050	0.100**	0.001	0.072**	0.106**	0.035
	(0.034)	(0.046)	(0.041)	(0.034)	(0.045)	(0.044)
Retirement product	0.045	0.147	0.027	0.038	0.173	0.015
	(0.105)	(0.160)	(0.123)	(0.108)	(0.164)	(0.122)
Home loan product	-0.038	0.028	-0.045	-0.044	0.142	-0.090
	(0.076)	(0.162)	(0.088)	(0.085)	(0.172)	(0.097)
Personal loan product	0.496***	0.184	0.595***	1.043***	1.120***	0.999***
	(0.105)	(0.231)	(0.121)	(0.101)	(0.228)	(0.123)
Savings product	0.198*	0.406**	0.072	-0.023	0.253	-0.170
	(0.120)	(0.161)	(0.154)	(0.133)	(0.189)	(0.168)
Term deposit product	0.277***	0.403***	0.182***	0.247***	0.369***	0.156***
	(0.040)	(0.062)	(0.050)	(0.040)	(0.062)	(0.050)
Transaction product	-0.324***	-0.437***	-0.230***	-0.209***	-0.376***	-0.071
	(0.039)	(0.063)	(0.054)	(0.043)	(0.065)	(0.057)
Home insurance policy	0.509***	0.736***	0.468***	0.349***	0.500***	0.334***
	(0.049)	(0.126)	(0.059)	(0.062)	(0.141)	(0.070)
Motor insurance policy	0.067	0.019	0.179**	0.105*	0.026	0.270***
	(0.058)	(0.077)	(0.087)	(0.058)	(0.078)	(0.095)
Constant	3.348***	3.060*	4.972***	3.455***	1.802	4.334***
	(0.212)	(1.579)	(0.403)	(0.245)	(1.644)	(0.427)
Observations	121,679	61,231	60,448	126,658	63,237	63,421
Customers	15,942	7,932	8,010	15,942	7,932	8,010
Data treatment for closed accounts	Missing	Missing	Missing	Zero	Zero	Zero
Sample	All	Age≤28	Age>28	All	Age≤28	Age>28
R-squared	0.0321	0.0461	0.0256	0.0271	0.0363	0.0302
Pred(Y): Non-Promotion: Control	$196.64	$158.84	$245.45	$192.67	$151.58	$246.66
Pred(Y): Non-Promotion: Treatment	$212.98	$164.56	$281.63	$211.79	$156.60	$294.06
Pred(Y): Promotion: Control	$177.20	$161.89	$195.12	$132.91	$130.44	$136.26
Pred(Y): Promotion: Treatment	$174.20	$158.90	$190.30	$130.99	$128.84	$133.13

Table 1.3: Effects of the treatment on monthly spend. Columns (1) – (3) show the results of panel data regression models for all customers as well as younger and older customers, with spend values for months after cancellation set to missing. Columns (4) – (6) show the same specifications with spend values for months after cancellation set to zero. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the bottom section of the table. The results indicate that customers exposed to the treatment spent more on their cards, that these effects were stronger for more experienced customers, and that these effects were attenuated by the promotion, which is consistent with H3, H5, and H7.

We believe that a more reliable test of whether tradeoff transparency helps customers make better choices is to analyze how they engage with their chosen products after acquiring them, which is our objective in the next two subsections.

3.2.3 Product usage (H3, H5, H7). In **Table 1.3**, Columns 1-3, we analyze the average monthly spending of all, less experienced, and more experienced customers. Average monthly spending captures the usage behavior of customers in all months during which they held an active credit card, including months when they had it, but didn't use it (e.g., monthly spend is $0), and excluding months after which they may have cancelled it (e.g., monthly spend is missing, because the customer no longer held the card). As such, monthly spend, as it is estimated in Columns (1-3), is a direct measure of contemporary customer engagement, in that it captures how the card was used while the card was held. Column (1) shows that the treatment led to a nominal, though insignificant, increase in average monthly spending (γ=0.080, p=0.122), that the promotion had no effect on spend (γ=-0.104, p=0.314), and that there was an insignificant interaction (γ=-0.097, p=0.111), when the effects were considered across all customers.

Similarly, Column (2) shows that the monthly spending of less experienced customers was neither influenced by the treatment (γ=0.035, p=0.675), nor by the promotion (γ=0.019, p=0.885), and that there was no interaction between the two (γ=-0.054, p=0.599). However, Column (3) illustrates that among customers who arrived during the non-promotion period, the monthly spending of more experienced customers was 14.74% higher among those who were provided with transparency into the offerings' tradeoffs (γ=0.138, p<0.05). Monthly spending was nominally lower among more experienced customers who arrived during the promotion period in the control condition (γ=-0.230, p=0.127), and there was a negative interaction, such that in aggregate, the treatment had no effect on monthly spending among more experienced customers during the promotion period (γ=-0.162, p<0.05).

Columns (4-6) present the same specifications, regressed on monthly spend where missing values due to account cancellations were filled with zero. As such, monthly spend, as it is estimated in Columns (4-6), is a direct measure of cumulative customer engagement, in that it captures how the card was used inclusive of the customer's choice of whether to retain it. Column (4) shows that during the non-promotion period, the treatment led to a marginal increase in spend (γ=0.095, p<0.10), increasing average monthly spending by 9.9% across all customers. It further shows that when cancellation dynamics are taken into account, the promotion decreased average monthly spend by 31.0% (γ=-0.371, p<0.01), and that there was a marginally significant interaction wherein the promotion attenuated the effects of the treatment to insignificance (γ=-0.109, p<0.10). Column (5) shows that the monthly spending of less experienced customers was neither influenced by the treatment (γ=0.033, p=0.722), nor by the promotion (γ=-0.150, p=0.229), and that there was no interaction between the two (γ=-0.045, p=0.684). However, Column (6) illustrates that during the non-promotion period, the monthly spending of more experienced customers was 19.2% higher among those who were provided with transparency into the offerings' tradeoffs (γ=0.176, p<0.01).

	(1) Pr(Late Payment)	(2) Pr(Late Payment)	(3) Pr(Late Payment)	(4) Pr(Late Payment)	(5) Pr(Late Payment)	(6) Pr(Late Payment)
Treatment	-0.117	-0.114	-0.135	-0.175**	-0.200*	-0.157
	(0.076)	(0.117)	(0.112)	(0.078)	(0.120)	(0.116)
Promotion	-0.471***	-0.295	-0.626***	-0.447***	-0.356	-0.512**
	(0.139)	(0.206)	(0.186)	(0.168)	(0.230)	(0.228)
Treatment x promotion	0.118	0.127	0.117	0.168	0.214	0.123
	(0.102)	(0.147)	(0.143)	(0.106)	(0.152)	(0.147)
Customer age	-0.074***	-0.170	0.027	-0.071***	-0.112	0.012
	(0.012)	(0.212)	(0.030)	(0.013)	(0.232)	(0.031)
Customer age²	0.001***	0.002	-0.000	0.001***	0.001	-0.000
	(0.000)	(0.005)	(0.000)	(0.000)	(0.005)	(0.000)
Customer tenure	-0.005***	-0.010***	0.000	-0.005***	-0.010***	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Customer tenure²	0.000***	0.000***	-0.000	0.000***	0.000***	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	-0.111**	-0.222***	0.008	-0.141***	-0.257***	-0.013
	(0.052)	(0.063)	(0.066)	(0.055)	(0.068)	(0.071)
Retirement product	-0.023	-0.077	0.022	-0.049	-0.118	0.004
	(0.095)	(0.116)	(0.150)	(0.101)	(0.123)	(0.161)
Home loan product	-0.752***	-1.430***	-0.603***	-0.752***	-1.383***	-0.617***
	(0.091)	(0.240)	(0.094)	(0.098)	(0.236)	(0.105)
Personal loan product	0.942***	1.062***	0.866***	0.901***	1.066***	0.792***
	(0.064)	(0.095)	(0.085)	(0.069)	(0.100)	(0.095)
Savings product	-0.772***	-0.937***	-0.624***	-0.775***	-0.904***	-0.658***
	(0.054)	(0.089)	(0.077)	(0.057)	(0.091)	(0.084)
Term deposit product	-0.749***	-1.340***	-0.468*	-0.793***	-1.414***	-0.493
	(0.235)	(0.473)	(0.280)	(0.259)	(0.511)	(0.301)
Transaction product	0.415**	1.485***	0.085	0.326*	1.336***	0.031
	(0.169)	(0.407)	(0.182)	(0.175)	(0.417)	(0.187)
Home insurance policy	0.105	0.308	-0.020	0.126	0.290	0.005
	(0.107)	(0.229)	(0.129)	(0.111)	(0.247)	(0.139)
Motor insurance policy	0.184	-0.014	0.213	0.241*	-0.103	0.332**
	(0.138)	(0.251)	(0.160)	(0.139)	(0.264)	(0.162)
Constant	-1.563***	-1.007	-3.792***	-1.483***	-1.573	-3.434***
	(0.295)	(2.459)	(0.648)	(0.325)	(2.711)	(0.695)
Observations	112,641	56,266	56,375	92,988	46,253	46,735
Customers	15,942	7,932	8,010	15,942	7,932	8,010
Account sample	All data	All data	All data	First six months	First six months	First six months
Sample	All	Age≤28	Age>28	All	Age≤28	Age>28
Wald Chi-Square	1177.32	970.85	577.18	1031.43	841.47	461.97
Pred(Y): Non-Promotion: Control	9.24%	9.68%	8.70%	8.82%	9.56%	7.96%
Pred(Y): Non-Promotion: Treatment	8.57%	9.02%	7.96%	7.87%	8.43%	7.15%
Pred(Y): Promotion: Control	6.78%	8.04%	5.66%	6.55%	7.62%	5.58%
Pred(Y): Promotion: Treatment	6.78%	8.11%	5.58%	6.52%	7.69%	5.44%

Table 1.4: Effects of the treatment on the probability a customer would make a late payment in a given month. Columns (1-3) show the results of panel data regression models for all customers as well as younger and older customers, regressed on the full sample. Columns (4-6) show the same specifications regressed on the first six months of observations for each customer, which balances the sample across customers acquired during the non-promotion and promotion periods. All models are estimated with random effects panel logistic regression, with robust standard errors, clustered by activation date. All models additionally include indicator variables for the week during which the customer activated his or her card, which are withheld from the table for parsimony. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the table's bottom section. The results indicate that the treatment had an insignificant or negative effect on the probability a customer would make a late payment. Column (4) demonstrates that the probability that a customer in the treatment condition, who was acquired during the non-promotion period was 10.8% less likely to make a late payment in any given month than customers acquired in the control condition. This result suggests the treatment has the capacity to improve the financial wellbeing of customers.

Monthly spending was 44.8% lower among more experienced customers brought in by the promotion in the control condition (γ=-0.593, p<0.01), and there was a negative interaction, such that in aggregate, the treatment had no effect on monthly spending among more experienced customers during the promotion period (γ=-0.199, p<0.05). With respect to monthly spending, taken together, these results suggest that although providing transparency to prospective customers can lead to greater engagement overall, which is consistent with H3, increases in engagement are most acute among more experienced customers, which is consistent with H5. More experienced customers who were provided transparency used their credit cards more intensively, spending 19.2% more on a monthly basis. The results further suggest that promotional efforts that provide incentives to encourage customers to choose one offering over another attenuate the effects of transparency on subsequent product usage, which is consistent with H7.

Importantly, although it may be reflective of firm profitability, we note that increased credit card spending, on its own, may not necessarily be emblematic of a better experience for customers. Indeed, higher spending rates could be detrimental to customers, if they meant customers were spending beyond their means. In a separate analysis, presented in **Table 1.4**, we analyze the probability that customers in different experimental conditions made late payments from month to month – the state of not having met the minimum required payment amount by the payment due date. Looking across the entire panel of data, in Columns (1-3), we find that the treatment had no effect on the probability a customer made a late payment – among all customers (γ=-0.117, p=0.124), less experienced customers (γ=-0.114, p=0.332), and more experienced customers (γ=-0.135, p=0.227). However, when in Columns (4-6) we consider the first six months of data for each customer – which balances the number of monthly observations among customers who opened accounts during the non-promotion (e.g., before November 1) and promotion (e.g., on or after November 1) periods – we observe that the probability of making a late payment was 10.8% lower among customers who were in the treatment condition (γ=-0.175, p<0.05), which suggests that the increased levels of engagement we observe do not come at the expense of customers' financial wellbeing. In fact, the evidence is supportive of the idea that tradeoff transparency can help customers make better choices, by selecting products they're more likely to use, and that will improve their financial wellbeing. In the next section, we examine the effects of transparency on customer retention, another important measure of customer engagement, that is both emblematic of customer experiences (e.g., customers who are having more favorable experiences with a service offering are more likely to continue using it), and firm performance (e.g., customers who are retained are generally more profitable than those who defect).

	(1) Pr(Retain6)	(2) Pr(Retain6)	(3) Pr(Retain6)	(4) Pr(Retain9)	(5) Pr(Retain9)	(6) Pr(Retain9)
Treatment	0.259	0.099	0.437*	0.249**	0.111	0.447***
	(0.202)	(0.340)	(0.231)	(0.115)	(0.183)	(0.163)
Promotion	-1.736***	-1.693***	-1.758***	-1.344***	-0.882**	-1.647***
	(0.256)	(0.633)	(0.275)	(0.231)	(0.422)	(0.294)
Treatment x promotion	-0.239	-0.055	-0.410	-0.223	0.114	-0.569**
	(0.220)	(0.365)	(0.258)	(0.165)	(0.258)	(0.226)
Customer age	-0.219***	-0.038	0.060	-0.094**	0.469	0.010
	(0.038)	(0.323)	(0.044)	(0.041)	(0.320)	(0.055)
Customer age²	0.003***	-0.002	-0.000	0.001***	-0.011	0.000
	(0.001)	(0.007)	(0.001)	(0.001)	(0.007)	(0.001)
Customer tenure	-0.009***	-0.011***	-0.009***	-0.005***	-0.007**	-0.006***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.003)	(0.002)
Customer tenure²	0.000***	0.000***	0.000***	0.000***	0.000*	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.046	0.011	0.065	0.078	0.039	0.118
	(0.074)	(0.110)	(0.095)	(0.091)	(0.100)	(0.133)
Retirement product	0.245	0.110	0.414*	-0.107	-0.255	0.036
	(0.154)	(0.221)	(0.229)	(0.155)	(0.176)	(0.235)
Home loan product	-0.595***	-0.775***	-0.539***	-0.564***	-0.545*	-0.538***
	(0.105)	(0.192)	(0.124)	(0.167)	(0.278)	(0.189)
Personal loan product	0.767***	0.408**	1.025***	0.246**	-0.007	0.497***
	(0.110)	(0.159)	(0.164)	(0.106)	(0.171)	(0.164)
Savings product	-0.188**	-0.141	-0.188**	-0.132	-0.070	-0.179
	(0.080)	(0.165)	(0.092)	(0.096)	(0.168)	(0.119)
Term deposit product	-0.617***	-0.574	-0.618***	-0.662***	-0.309	-0.796***
	(0.167)	(0.352)	(0.203)	(0.235)	(0.484)	(0.294)
Transaction product	1.603***	2.235***	1.312***	1.402***	2.034***	1.074***
	(0.125)	(0.269)	(0.159)	(0.222)	(0.413)	(0.282)
Home insurance policy	0.084	0.504	-0.012	-0.068	0.777	-0.233
	(0.154)	(0.383)	(0.162)	(0.170)	(0.668)	(0.199)
Motor insurance policy	0.082	0.585	-0.045	0.027	0.744	-0.109
	(0.198)	(0.492)	(0.224)	(0.202)	(0.561)	(0.235)
Constant	7.509***	5.895	1.963**	3.739***	-2.343	1.345
	(0.711)	(4.006)	(0.914)	(0.773)	(3.835)	(1.219)
Observations	15,942	7,864	8,010	6,314	3,231	3,083
Customers	All	Age≤28	Age>28	All	Age≤28	Age>28
Pseudo R2	0.0845	0.0916	0.0977	0.0551	0.0509	0.0948
Pred(Y): Non-Promotion: Control	98.07%	98.32%	97.76%	93.40%	92.99%	93.38%
Pred(Y): Non-Promotion: Treatment	98.50%	98.47%	98.53%	94.75%	93.66%	95.61%
Pred(Y): Promotion: Control	90.50%	92.14%	88.99%	79.43%	84.98%	75.00%
Pred(Y): Promotion: Treatment	90.67%	92.43%	89.23%	79.83%	87.54%	72.87%

Table 1.5: Effects of the treatment on customer retention rates. Columns (1) – (3) and columns (4) – (6) show the results of cross-sectional data regression models of all, younger, and older customers for retention six and nine months after card activation, respectively. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the bottom section of the table. The results indicate that transparency can lead to increased customer retention, which is consistent with H4, that the effects are strongest for more experienced customers, which is consistent with H6, and that the effects are dampened during the promotion, which is consistent with H8. Marginal effect estimates for retention after six and nine months are provided for each experimental condition.

3.2.4 Retention (H4, H6, H8). In **Table 1.5**, Columns (1-3) and (4-6) show the effects of tradeoff transparency on the retention of all customers, less experienced customers, and more experienced customers, after six and nine months, respectively. Columns (1) and (4) provide evidence that, looking across all customers, the treatment had an insignificant effect on retention after six months (δ=0.259, p=0.198), but that it increased retention after nine months (δ=0.249, p<0.05). After nine months, cancellations among customers in the treatment group were 20.5% lower, which is consistent with H4, and provides evidence that revealing the hidden tradeoffs to prospective customers can improve their decision making in ways that have positive implications for the long run trajectories of service relationships.

Interestingly, Columns (2) and (5) reveal that, consistent with the prior analysis on spend, there's no evidence that the positive implications of the treatment for retention accrued to customers with less category experience. Less experienced customers in the treatment group were no more nor less likely to be retained after six (δ=0.099, p=0.771) or nine months (δ=0.111, p=0.546). In contrast, Columns (3) and (6) show that more experienced customers who received the transparency treatment were more likely to have been retained after six (δ=0.437, p<0.10) and nine months (δ=0.447, p<0.01), which is consistent with H6.

The positive effect of tradeoff transparency on product retention was most acute for customers with prior experience with the service category. Among these more experienced customers, nine months into their service relationships, cancellation rates were 33.7% lower for those who had been randomly assigned to experience tradeoff transparency before making their product choice (falling from 6.62% to 4.39%). Finally, Column (6) demonstrates an interaction, wherein the effect of the treatment on retention among more experienced customers was negatively moderated after nine months for customers who were attracted by a promotion (δ=-0.569, p<0.05). Similarly, the negative coefficients on the interaction terms reduce the positive and significant effects of the treatment on retention to insignificance during the promotion period ($\chi^2 s$<0.53, ps>0.469). These results are consistent with H8, in that they demonstrate how the positive effect of tradeoff transparency on product retention can be diminished among customers attracted to an offering by a promotion.

4. General discussion

Through a large-scale field experiment conducted with 393,036 customers of a nationwide retail bank, we have investigated the impact of providing prospective customers with transparency into a service offering's tradeoffs on rates of acquisition, product choice, and subsequent long-term engagement. Our results suggest that although providing prospective customers with tradeoff transparency has no net effect on overall rates of acquisition, it can affect customers' choices – causing them to select different offerings than they would in the absence of tradeoff transparency. Moreover, we find evidence that the information tradeoff transparency provides can help customers make better choices, leading to long-term outcomes that are improved for customers and firms alike. Customers who were provided tradeoff transparency and chose to move forward by opening an account exhibited higher levels of engagement during their first nine months of service, effects that were especially pronounced among customers who had prior experience with the service category. Monthly spending was 9.9% higher, and after nine months, cancellation rates were 20.5% lower, among all customers who received the tradeoff transparency treatment. For more experienced customers, monthly spending was 19.2% higher, and rates of cancellation after nine months were 33.7% lower among those who received the tradeoff transparency treatment. Importantly, we find that increased engagement levels do not come at the expense of customers' financial wellbeing. During the initial six months of their relationships, customers who experienced tradeoff transparency were 10.8% less likely to make late payments on a monthly basis. These results provide evidence that being more transparent with customers – not just about an offering's strengths, but also about its tradeoffs – can help customers make choices that lead to better long-run outcomes for customers and firms alike. Finally, we find that the effects of transparency on engagement and retention are attenuated in the presence of a promotion. The results are consistent with the idea that customers who organically seek out an offering are more likely to benefit from tradeoff transparency than those who are attracted or influenced by a promotion.

We note that the results documented in this paper arose from a field experiment conducted in a highly-regulated service environment, where providers are legally required to fully disclose the terms and conditions of their offerings to prospective customers. In this context, tradeoff transparency merely increased the salience of information that was already being provided, and its effects were substantial. It stands to reason that the effects of tradeoff transparency may be even stronger in less-regulated contexts, where such disclosures are not required and customers face even greater information asymmetries.

4.1 Managerial implications

Taken together, these results have important implications for practice, while also raising additional questions that will serve as fruitful opportunities for future research.

4.1.1 Tradeoff transparency may not harm acquisition rates. Although research has shown how the revelation of unflattering information can harm demand, and conventional wisdom has long suggested the fallacy of marketing an offering's weaknesses, our results demonstrate that voluntarily providing customers with transparency into an offering's tradeoffs may not harm rates of acquisition. These results are important, because they suggest that the costs of being fully open with prospective customers that many managers perceive (e.g., that prospective customers who are shown the tradeoffs in an offering may seek service elsewhere) may be misplaced.

4.1.2 Tradeoff transparency may lead to better long-run outcomes for customers and firms. Although the provision of transparency in our study had an insignificant effect on overall rates of customer acquisition, we find that it caused customers to make different product choices than they would make in the absence of transparency. Specifically, we observed that customers in different segments, who were given the same tradeoff information, selected different credit cards. This pattern of results is consistent with the idea that the revelation of information about a product's tradeoffs can help shape their decisions. Furthermore, we find that transparency improves customer engagement – both by increasing product usage and rates of retention over the longer-term. Customers who were randomly-selected to experience transparency into each offering's tradeoffs, and who moved forward in opening an account, used their cards more intensively, spending 9.9% more per month, and were 10.8% less likely to make late payments on a monthly basis.

Furthermore, customers who experienced transparency were 20.5% less likely to cancel their credit cards during the first nine months of their relationships. To the extent that heightened engagement and retention are behavioral indicators of a better customer experience, and a more profitable service relationship, these results suggest that providing transparency to prospective customers may be mutually beneficial.

4.1.3 There may be a tradeoff between acquisition and retention-based strategies. Midway through the experiment, our research partner launched a promotion in which they offered a financial incentive to prospective customers who opened a credit card in the Low Rate family, and who spent $1,000 during the first three months of their service relationships. From an acquisition perspective, the efficacy of this promotional strategy is evident. Traffic to the credit card website increased, and the conversion rate of browsers to buyers was enhanced, such that the number of new credit card accounts opened per day went up by 47.3% during the promotion period. Furthermore, the probability a customer would open a promoted credit card went up – during the promotion the probability of opening a low rate or low rate gold credit card increased by 28.1% and 18.0%, respectively – highlighting the promotion's capacity to increase sales and influence customer choices. However, customers acquired during the promotion spent a third less on a monthly basis, due in part to the fact that they were 3.8 times more likely to cancel their cards during the initial months of their relationships than those who were not acquired through a promotion. These results highlight a potential tradeoff for managers between marketing-based promotional strategies that are optimized around increasing rates of customer acquisition, and operations-based transparency strategies that are optimized around increasing customer compatibility and long-term engagement. Future research could delve more deeply into this tradeoff – for example, by examining whether over a longer timeframe, the gains in engagement and retention brought about by tradeoff transparency might outweigh the improvement in acquisition rates promotions can foster.

4.1.4 Tradeoff transparency is especially effective for customers who have prior category experience. Consistent with prior research, our results provide evidence that transparency may be particularly effective for improving relationship quality and engagement among more experienced customers. In our experiment, more experienced customers (operationalized as customers older than 28 years of age), who received the transparency treatment, used the credit cards they chose more intensively, exhibiting cumulative spending that was 19.2% higher after nine months, than those who did not experience transparency. They also were 33.7% less likely to cancel their credit cards during the first nine months of their relationship. These results reveal the incredible promise of tradeoff transparency, for customers and companies alike, in helping customers make more well-informed decisions. However, it's interesting that less experienced customers did not exhibit similar gains in subsequent engagement and retention. Less experienced customers (those 28 years of age and younger), who were treated with tradeoff transparency, spent 3.3% more during a typical month, and were 9.6% less likely to cancel their accounts after nine months – results that are statistically indistinguishable from zero. We note, however, that although the benefits of tradeoff transparency for relationship quality and engagement are especially pronounced among more experienced customers, supplementary interaction models presented in full in the online appendix do not find age (our proxy for experience) to significantly moderate the effects of tradeoff transparency on spending (*ps*>0.16) or retention (*ps*>0.21). Nevertheless, our results highlight an opportunity to explore strategies for helping to better inform less-experienced prospective customers who do not significantly benefit from tradeoff transparency in our analyses.

4.1.5 Promotions may crowd-out the benefits of providing prospective customers with tradeoff transparency. Finally, our results suggest that the effects of providing prospective customers with transparency into an offering's tradeoffs may be attenuated by promotions, which themselves are designed to influence customer choices. More experienced customers, in particular, who exhibited the greatest increases in product usage and retention in response to tradeoff transparency, also exhibited an interaction effect, wherein those gains were mitigated, though not reversed, in the presence of a promotion. Importantly, in our experiment, the promotion in question offered financial incentives for only one family of credit cards, which may have created a conflict for some consumers between the card that offered a bonus and the card that best fit their needs and preferences.

Future research could more holistically disentangle whether promotions and tradeoff transparency are necessarily in conflict by treating every offering with the same promotion. Such a design would afford a better understanding of the mechanisms that drive the crowding-out effect we observe – whether it's driven by the conflict between fit and financial incentives some customers may have experienced, or whether it's driven by promotions attracting customers who are less sensitive to the effects of tradeoff transparency. Owing to the demonstrated efficacy of promotions to attract new customers, and tradeoff transparency to foster decision making that boosts long-run engagement, innovations that achieve the best of both worlds could enhance customer experiences while unlocking considerable value for firms.

4.2 Limitations and opportunities for future research

An important caveat of this work is that since the focal firm offered a portfolio of products, many of the tradeoffs presented for particular cards could be overcome by choosing a different card from the bank. Indeed, as shown in **Figure 1.2**, some of the tradeoffs presented specifically compare service attributes with those of other cards, such as "the annual fee of $59 is higher than our Low Fee credit card," or "the purchase interest rate of 20.24% p.a. is higher than the Low Rate and Low Fee credit cards." It is possible that one of the reasons acquisition rates didn't fall in this study was that tradeoff transparency helped facilitate comparisons among the bank's own offerings. Future research could examine the effects of tradeoff transparency in contexts where particular tradeoffs cannot be overcome by selecting another offering from the firm. Relatedly, future research might investigate how tradeoff transparency that varies in valence has differential effects on consumers. In the present research, benefits and tradeoffs were quite evenly-balanced, but consumers may draw fundamentally different inferences as the ratio of good to bad changes. Indeed, it may be the case that companies need to improve the quality of their underlying offerings to a particular level before being genuinely transparent about tradeoffs becomes a viable strategy.

These results are consistent with prior research that has shown how the act of being voluntarily transparent with information that customers perceive to be sensitive can engender higher levels of trust in the firm and attraction to its offerings (Mohan et al. 2020). Indeed, in an online experiment conducted with

201 participants (62.2% male, M=39.18 years old), and presented in full in the online appendix, we find evidence that tradeoff transparency can increase trust and brand preference. 71% of participants ($p<0.01$, one-sided) presented with credit card webpages from competing banks reported the bank that provided tradeoff transparency to be more trustworthy, and 69% ($p<0.01$, one-sided) reported being more likely to apply for a credit card with the transparent bank. An opportunity for future research would be to delve more deeply into the mechanisms that underlie these effects, exploring the conditions under which tradeoff transparency is more and less likely to engender trust, and how it affects the relationship with the firm more holistically (e.g., does it affect peoples' willingness to engage with other offerings).

Likewise, future research could explore whether the provision of tradeoff transparency is generally appealing, or if there are individual differences among appetites for transparency, such that transparent organizations attract different types of customers. In the online appendix, we present evidence that customers who chose to open credit cards in the treatment and control conditions of our experiment were substantively similar on observable characteristics. These results suggest that in this setting, improved engagement arose from tradeoff transparency's capacity to help otherwise-similar customers make choices that were better aligned with their needs and preferences, rather than by attracting fundamentally different customers. We note, however, that results may differ in other settings – particularly when the marketed offering primarily targets new customers, rather than extends an ongoing service relationship, as may be the case with bank-issued credit cards. Similarly, research on the psychology of choice has demonstrated how providing more information that complicates a choice can lead a customer to opt out of choosing (Gourville and Soman 2005, Iyengar and Lepper 2000). Future research could investigate how tradeoffs could be most effectively communicated to help customers make better-informed decisions as easily as possible.

4.3 Conclusion

Conventional wisdom and common practice dictate that service firms should emphasize the advantages of their offerings and downplay the tradeoffs when marketing to prospective customers. We suggest that taking a different approach – providing transparency into both an offering's advantages and its tradeoffs, in order to help customers make more well-informed decisions – can lead to better outcomes over the long run for everyone involved. We hope that the present work will lead to more research in this area, and influence practice, in order to foster better customer experiences, and more engaging and profitable service relationships among customers and the organizations that serve them.

An Empirical Study of Time Allotment and Delays in E-commerce Delivery

Maya Balakrishnan, MoonSoo Choi, Natalie Epstein
Harvard Business School, Technology & Operations Management Unit
mbalakrishnan@hbs.edu, mochoi@hbs.edu, nepstein@hbs.edu

Problem definition: We explore the relationship between time allotment and delivery outcomes in e-commerce delivery. Furthermore, we seek to identify relevant features for predicting order delays and study how various facets of real-time delivery data can be used to improve prediction accuracy.

Methodology: We use the JD.com transaction dataset provided by Shen et al. (2020). We first employ a Regression Discontinuity design to examine the effect of exogenous variations in time allotment between same-day and next-day orders on delivery outcomes, such as delays and duration of the process. Subsequently, we fit random forest classification models to predict delays and identify the key predictor variables. Finally, we construct and incorporate load-related features in our prediction models to explore its impact on the overall model performance.

Relevance: We draw methods from causal inference (to understand the relationship between allotted time and delivery outcomes), machine learning methods (to predict future delivery delays), and fluid model in queueing (to construct load-related features) to identify sources of operational inefficiency.

Results: We find that (*i*) increasing the allotted time for an order increases the duration of its delivery process, (*ii*) more allotted time reduces the likelihood of an order being late or delayed, (*iii*) adding information from early parts of the delivery significantly increases accuracy when predicting future delays, and (*iv*) load-related features can boost the performance of delay prediction models, but not to the same extent as information on earlier legs of delivery.

Managerial Implications: Time allotment comes with an inherent tradeoff between chances of delay and delivery duration (and potential implications for customer satisfaction), so managers need to carefully evaluate such tradeoff. For developing prediction models, managers can use order characteristics and information about earlier stages of the delivery process to predict delays before they occur and take actions to mitigate such delays.

1. Introduction

As today's e-Commerce is becoming increasingly complex and customers expect timely deliveries, retailers have to coordinate their order fulfillment processes seamlessly to deliver their orders on time (Mao et al. 2019). For retailers, continuous improvement in logistic processes is indispensable in acquiring and retaining customers, as delivery outcomes are closely related to profitability and customer loyalty (Fisher et al. 2019). Although the order fulfillment process itself does not directly impact customers' perceptions of the products they receive, they nonetheless play a pivotal role in ensuring that customers are satisfied with the experience of placing and receiving online orders. A recently emerging body of work shows that both the promised delivery time, as well as the actual amount of time an order takes to be delivered, are critical elements of logistic quality that impact sales and returns (Yu et al. (2020), Cui et al. (2020)). Therefore, online retailers need to consider how they can improve their logistics processes to promise rapid deliveries and, simultaneously, ensure that orders arrive on time.

Within the context of online retail logistics, we study the relationship between the time allotment, i.e., the amount of time available to deliver an order, and delivery performance outcomes, such as the probability of delay and lateness, the duration of the delivery, and the average action time of the delivery. Furthermore, we aim to predict delays based on order characteristics and identify the key features that help in predicting these delays. Finally, we examine whether adding information about in-process delivery — such as duration of delivery in each stage or warehouse workloads — increases the accuracy of our prediction models.

To examine these questions, we use a large-scale dataset from JD.com, provided by the company and the 2020 M&SOM Data-Driven Challenge (Shen et al. 2020). To study the effect of time allotted on the delivery process outcomes, we use a Regression Discontinuity (RD) design. We exploit a time of day cutoff that determines whether an order with a 1-day promise will be delivered within the same day or the next day, as a source of exogenous variation of allotted time to deliver an order. Specifically, orders placed before 11 am have a same-day promise, whereas orders placed after 11 am but before 11 pm are promised to be delivered before 3 pm on the next day.

For predicting delays, we fit a Random Forest classification model to classify whether an order will be delayed or on time. We use this model to identify features of an order that are highly predictive of whether the order will be delayed. Moreover, we augment this baseline model with information about the early stages of the delivery process. The main idea is to examine whether including such real-time information about the early stages in the model significantly improves prediction accuracy. Finally, we use fluid models in queues developed in (Hall 1991, Newell 2013) to construct features characterizing workload for each warehouse at each hour. We then incorporate these load-related features in our prediction model and evaluate its performance.

Overall, we find that (*i*) increasing the allotted time to deliver an order decreases its probability of delay. However, we also find that (*ii*) increasing the allotted time for delivery disproportionately increases each leg's duration in the delivery process, relative to the increased amount of time available. Additionally, we find that (*iii*) our delay prediction models exhibit accuracy levels higher than those of naïve classifier model. Moreover, we find that (*iv*) adding information about earlier stages of the delivery process significantly increases our models' accuracy levels. Finally, results suggest that (*v*) including load-related features can also improve prediction performance, but not to the same extent as including real-time data on earlier stages of the delivery. Taken together, we believe that these results can inform eCommerce companies how resources should be allocated for their delivery systems to ensure orders are delivered on time. Furthermore, our prediction models can help managers identify early on which orders are at risk of arriving past the promise time, allowing them to reprioritize resources to help them stay on track.

The remainder of the paper is organized as follows. Section 2 summarizes relevant literature related to this work. In Section 3, we present the empirical setting and dataset description. In Section 4, we describe the Regression Discontinuity models and present the results, and in Section 5, we present the Machine Learning models and their results. Augmenting Section 5, Section 6 leverages fluid model to develop and incorporate load-related features. Section 7 discusses managerial implications and concludes.

2. Literature Review

2.1. Delivery and customer satisfaction

In online retail, reliability in fulfillment and delivery processes is key to increasing and maintaining customer satisfaction. For instance, Rao et al. (2014) show that when the delivery timeline is reliable, products are less likely to be returned. This indicates that increasing the consistency with which orders arrive on time could decrease returns and increase profits. Furthermore, they show that this effect is stronger for orders that are promised to be delivered quickly, which is of particular relevance to our analysis, since we focus on orders promised to be delivered within the same or next day.

In this same context, Fisher et al. (2019) examine how the duration of the delivery process affects profits. They demonstrate that in the case of an online retailer, when the physical delivery speed increased, total revenue and net profit both increased significantly. Cui et al. (2020) show that more aggressive promises (meaning promising shorter delivery times to customers) can increase sales and profits. However, this strategy can also harm the retailer and cause an increase in returns, if the more aggressive promises aren't actually accompanied by an increase in delivery speed, leading to more delays. Such disadvantage is evident in the field experiment in Yu et al. (2020), and as implied by the rational abandonment model in Mandelbaum and Shimkin (2000), where customers who were given more optimistic estimates but were not served on time exhibit higher likelihood of abandonment.

In addition to speed and consistency of the delivery process affecting sales and returns, in a similar setting to the one we explore, Cui et al. (2019) found that when a high quality delivery option was removed from Alibaba, sales dropped by 14.56%. And Bray (2019) shows that for online retail orders from Alibaba, customers punish idleness later in the delivery process more than idleness occurring early on in the process. Bray (2019) further finds that pushing the majority of delivery process actions closer to the time of delivery can provide a comparable increase in customer satisfaction to decreasing the overall duration of the delivery process.

While we do not directly study customer satisfaction with the delivery process, we contribute to this literature by showing how the time allotted to a delivery can affect these metrics of the delivery process (such as delay probabilities and durations), which are in turn drivers of customer satisfaction.

2.2. Parkinson's Law

Our work also relates to Parkinson's Law, which states that: "work expands so as to fill the time available for its completion" (Parkinson and Osborn 1957), which has numerous implications for project management and logistics (Gutierrez and Kouvelis 1991). Prior work by Hasija et al. (2008) show how workers have incentives to slow down their work when capacity exceeds demand. Marin et al. (2007) study queues in an airport and propose that server behavior should also be considered when designing a system, as workers might change the speed of the system depending of the length of the queue. Kc and Terwiesch (2012) show that in medical settings, the speed of service can change depending on the workload. In our work, we show that increasing the allotted time to deliver an order is related with actually taking more time to complete the delivery, consistent with Parkinson's law. We also find that this increase in delivery times due to increased allotted time is not necessarily at the expense of increasing delays. This is especially important when we consider work from Lu et al. (2013) and Musalem et al. (2017), who find that (a) customers are sensitive to uncertainty about service quality, and (b) customers display an asymmetric response such that they are more sensitive to delays than early deliveries.

2.3. Prediction models

Finally, by fitting prediction models and exploring their value in the retail industry, our research seeks to add to the increasing body of retail analytics literature. Just to name a few, Caro and Gallien (2012), Cui et al. (2018), Ferreira et al. (2016), and Perakis et al. (2018) have developed and implemented data-driven decision support tools in abetting retail companies with predicting demand and determining optimal pricing schemes. With respect to predicting delay and task duration, Guo et al. (2019) has developed a prediction system with regression trees that could be run in real-time to identify any late passengers for their connecting flights. Similarly, the tree-based prediction model that we develop also shows the value of using real-time duration information to improve the accuracy of detecting any potential delays early.

3. Setting and Data Description

This section is organized as follows. First, we present the specific empirical setting and describe the dataset in detail used to conduct our analyses. Second, we describe the delivery process and associated variables. Finally, we present summary statistics of the data.

3.1. Empirical Setting

We conduct our analyses using data provided by JD.com as part of the 2020 M&SOM Data Driven Research Challenge (Shen et al. 2020). JD.com is China's largest online retailer, with annual revenues of $67.198B USD in 2018. In addition to offering a vast selection of products across all major categories, they own and operate their own logistics network that ships out their own products and provides third party shipping and warehousing for other brands. JD.com prides themselves on the quality of their products as well as the speed and scale of their logistics network. Their website (JD.com 2020) describes their logistics network saying:

> "JD.com has one of the largest fulfillment infrastructure of any e-commerce company in the world. Currently, JD.com operated 28 "Asia No. 1" logistics parks, which are among the largest and most automated smart fulfillment centers in Asia... JD.com is able to achieve rates of approximately 90% of orders delivered the same or next day, a rate of fulfillment that no other e-commerce company of JD.com's scale can match globally."

Given the importance of rapid delivery to JD.com's business, we choose to focus our analyses on the portion of data that corresponds to orders that were promised to be delivered within the same or next day.

As part of the M&SOM Data Driven Research Challenge, JD.com provided us with a snapshot of their data from March 2018 in a single product category. This data includes browsing and purchasing behavior, shipping and delivery data, as well as customer demographics information. For our analyses, we use the following tables:

- Orders table: contains a single row for each unique order-SKU pair. If a customer placed an order on March 3 with 3 different SKUs in the order, there would be 3 corresponding rows in this table. Each row also contains information at the order level, such as when the order was placed and which warehouse it was shipped from, as well as SKU-level information such as the quantity of the SKU ordered and its price. Of particular importance is the "promise" field, which states JD.com's expected number of days in which the order will be delivered. A promise of 1 corresponds to same- and next-day delivery.
- Delivery table: contains a single row for each unique package. A single order can be shipped out in multiple packages, but the vast majority of orders (99.98%) are shipped out and sent to the customer in a single package. Given that we conduct our analyses on an order level, we use the delivery timestamps for the package in that order that is last to be delivered to the customer. As part of the delivery information, we have timestamps on when the package was shipped out from the warehouse, when it arrived at the delivery station, and when it was delivered to the customer.
- Users table: contains a single row for each unique JD.com user who purchased an item from the product category of interest in March of 2018. Each row contains information on a user such as how long they have been a member of JD.com, and whether they are a PLUS (premium) member, as well as information that is imputed by JD.com, such as the user's age group and marital status.

In the Online Appendix A.3.1 we present the steps to construct the subset of data that we use for our analyses. The result of this process is a single table containing 127,054 unique orders with delivery and customer information.

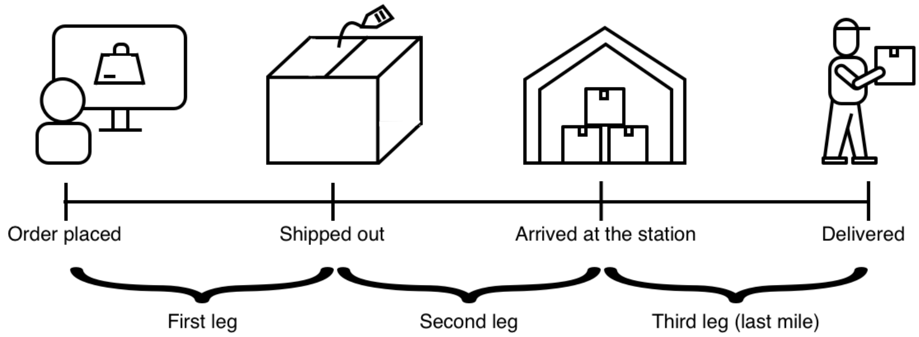

Figure 2.1 Delivery Process

3.2. Delivery process and variables

Using the features in this resulting table, we can construct the delivery process. In Figure 2.1 we show a scheme of the delivery process, which is comprised of 4 timestamps: (1) "Order placed" is the time of the day when the order was placed by the user, to the second level, (2) "Shipped out" is the time when the order ships from the warehouse, to the hour level, (3) "Arrived at station" is the time, to the hour level, when the order arrives at the last-mile station, and (4) "Delivered" is the time, to the second level, when the order arrives at the customer's place.

Given that for each order we have four timestamps as defined above, we can represent each order's journey as a series of 3 legs. We define the first leg as the time between when the order is placed up to when it is shipped out. Similarly, we define the second leg as the time between when the order is shipped out and when it arrives at the last-mile station. Finally, we define the third leg as the time from when the order arrives at the last-mile station to when it is delivered.

Within our dataset, we define additional variables that will help us construct our models:

- Fractional Order Time: For orders with a promise of 1, the time of day when they are placed determines whether they will be delivered on the same day or the next day. Therefore we define Fractional Order Time as the time of day, in fractional hours, when the order is placed, ignoring the date. Fractional Order Time is a number ranging from 0 to just under 24. For example, if an order is placed at 3:45 pm, the Fractional Order Time for that order will be 15.75.

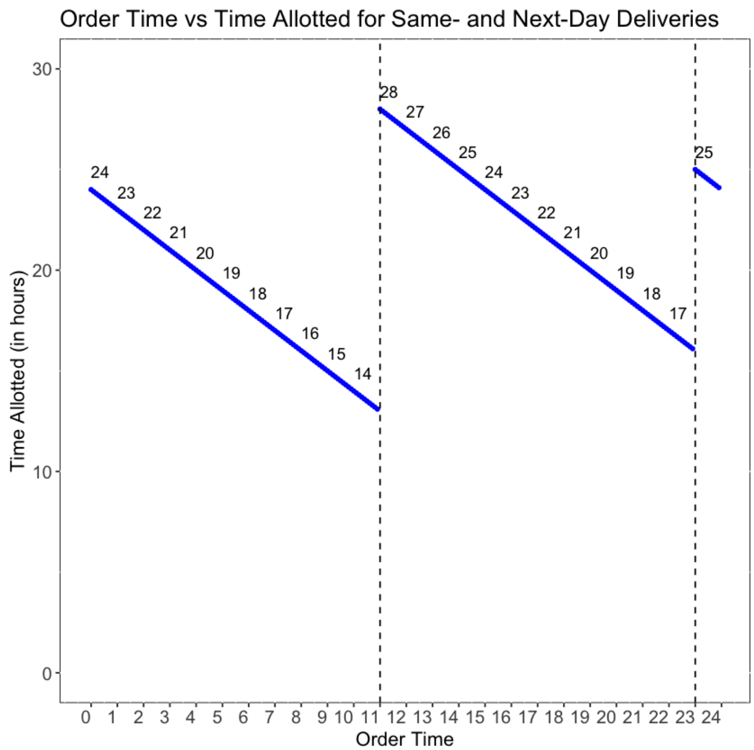

Figure 2.2 Order Time vs. Time Allotted for same- and next-day promise.

- Promise Time: JD.com's policy states, "When promise = 1, this refers to the standard same- and next-day delivery promise: Orders placed before 11 am will be delivered on the same day, and orders placed before 11 pm will be delivered before 3 pm on the following day." Promise Time represents the latest possible time at which an order can arrive at the customer's house, according to this delivery policy.
- Time Allotted: This represents how many hours an order has to be delivered within. It is calculated by taking the difference between Order Time and Promise Time and converting that to a fractional number of hours. We can see this in Figure 2.2.

3.3. Summary statistics

Our final data sample is comprised of 127,054 unique orders which are promised to be delivered within the same or next day. Of these orders, 28,574 (22.49%) are delayed, meaning they arrived at the customer's house after they were promised to be delivered. For these delayed orders, the mean time duration for the order to be fulfilled is 41.17 hours, with a standard deviation of 28.02 hours. The mean first leg duration for the delayed orders is 16.62 hours, the mean second leg duration is 13.47 hours, and the mean third leg duration is 11.08 hours. On average, orders that are delayed arrive 21.71 hours after they were promised to arrive. See Table A2.1 in the Online Appendix for details on delayed orders.

For the 98,480 (77.51%) of orders that arrive on time, their mean duration is 16.10 hours, with a standard deviation of 5.35 hours. For these on-time orders, their mean first leg duration is 2.07 hours; their mean second leg duration is 10.04 hours; and their mean third leg duration is 3.98 hours. On average, on-time orders arrive 4.85 hours early (before they were promised to be delivered). See Table A2.2 and Table A2.3 in the Online Appendix for details on on-time orders and a summary of the duration for on-time orders' first legs, separated by the hour of the day when they were placed, respectively.

Table A2.4 in the Online Appendix presents further summary statistics across all the orders in our sample, looking at features of each order's content and value.

Orders are placed throughout the course of each day as well as throughout the month that our dataset captures. Figure 2.3 shows the distribution of ordering behavior over the course of the day. Although we do not consider variation in the day of the month when orders are placed as part of our analyses, we do take advantage of the distribution in time of day when orders are placed.

4. Regression Discontinuity

This section is organized as follows: (1) we present the discontinuity that allows us to use an RD framework to study how time allotted affects delivery outcomes, (2) we define the outcomes of interest, (3) we present the model specification for each outcome, (4) we validate the necessary assumptions for RD, and (5) we present the results.

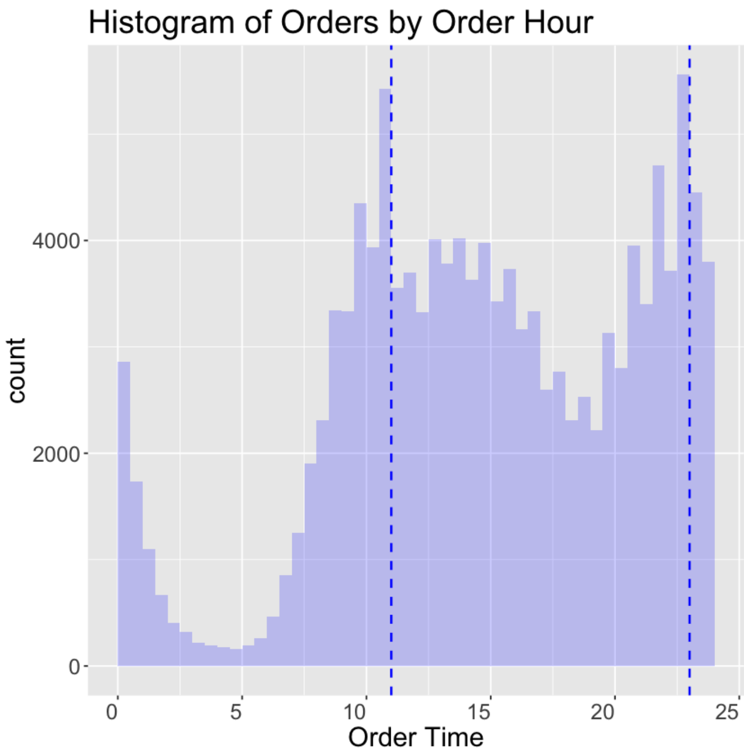

Figure 2.3 **Number of orders by hour of the day. The vertical line represents the 11:00 and 23:00 cutoff.**

4.1. Discontinuity

We use an exogenous variation in time allotted driven by an 11 am cutoff that determines whether an order will be delivered within the same day or the next day. This allows us to conduct a regression discontinuity (RD) analysis on how the time allotted affects various delivery outcomes. Given that there are no substantive differences between the orders placed at 10:59 am and 11:01 am (e.g., in terms of order, delivery, and customer characteristics) – other than the fact that orders placed just after 11 am allow themselves an extra 15 hours to be delivered compared to orders placed just before 11 am – regression discontinuity allows us to capture the effect that the extra allotted time has on the delivery outcomes of orders (Lee and Lemieux 2010). As described in the prompt in Shen et al. (2020), "for orders with promise = 1, Orders placed before 11 am will be delivered on the same day, and orders placed before 11 pm will be delivered before 3 pm on the following day.". A visual representation of how the time allotted varies with the time of day when an order is placed can be seen in Figure 2.2.

4.2. Outcomes

Based on the dataset presented in Section 3 we construct several outcome variables to understand how time allotted affects different aspects of the delivery process.

(i) Delayed: An order is Delayed if it arrives at the customer's house after it was promised to arrive. Delayed is a binary variable which is set to 1 if the order's Arrival Time occurs after the Promise Time, and 0 if the Arrival Time is before the Promise Time.

(ii) Late in Leg t ($t = 1, 2, 3$): An order is late in leg t if the duration of the leg is longer than the duration of most of the other on-time orders that were placed at the same time of day. Given that the Time Allotted for an order is driven by the time of day the order is placed, we define Late Leg t as a binary variable which is equal to 1 if the order's t'th leg took longer than the 95th percentile of all on-time orders' t'th legs that were placed within the same hour of the day, and 0 otherwise.

(iii) Duration of Leg t ($t = 1, 2, 3$): It is the duration, converted to fractional hours, of leg t, as illustrated in Figure 2.1.

(iv) Ratio for Leg t ($t = 1, 2, 3$): Because there may be inherent differences among the three legs in their duration and how they process their orders, we define the duration Ratio as $\frac{\text{Duration leg t}}{}$ Time Allotted to normalize the duration of each individual leg.

(v) Average Action Time: The action time is a fractional measure between 0 and 1, defined as the cumulative duration of all legs thus far divided by the total duration of the delivery. We only focus on action times of legs 1 and 2, since the third leg's action time will always be equal to 1.00. The "average action time" is defined as the average of all intermediate action times that comprise the order's delivery process. For our dataset, each order's average action time is calculated as follows:

$$\frac{1}{2}\left(\frac{\text{Ship out time} - \text{Order time}}{\text{Duration}} + \frac{\text{Station arrival time} - \text{Order time}}{\text{Duration}}\right)$$

4.3. Specification

To set up the RD framework, we define the variable AFTER_11$_i$ for order i as follows:

$$\text{AFTER\_11}_i = \begin{cases} 1 \text{ if Order\_Frac\_Hour} \geq 11 \\ \\ 0 \text{ otherwise} \end{cases}$$

The variable Order_Frac_Hour represents the "fractional order time" defined in Section 3. We also define the running variable τ_i, which represents the distance (in hours) of the time the order is placed from the 11 am threshold, which is numerically equivalent to $\tau_i = \Big(\text{Order\_Frac\_Hour}_i - 11\Big)$.

Additionally, we define X_i as a set of other covariates associated with an order, representing customer demographics (e.g., age group, education) and order characteristics (e.g., coupon usage, order cost).

(i) Delayed and (ii) Late in Leg t:

For estimating the effect of increased allotted time on delay probability, we have:

$$P(\text{Delay})_i = f\Big(\alpha + \beta \cdot \tau_i + \gamma \cdot \text{AFTER\_11}_i + \delta \cdot \tau_i \cdot \text{AFTER\_11}_i + \zeta \cdot X_i + \epsilon_i\Big)$$

where $P(\text{Delay})_i$ is the probability of order i arriving at the customer's house after it was promised.

For estimating the effect of increased allotted time on the probability of each leg being late, we have:

$$P(\text{Late\_leg}_t)_i = f\Big(\alpha_t + \beta_t \cdot \tau_i + \gamma_t \cdot \text{AFTER\_11}_i + \delta_t \cdot \tau_i \cdot \text{AFTER\_11}_i + \zeta_t \cdot X_i + \epsilon_{i,t}\Big)$$

where $P(\text{Late\_leg}_t)_i$ is the probability of leg t ($t = 1, 2, 3$) for order i taking longer than 95% of all on-time orders (placed at the same hour of the day).

The functional form $f(\cdot)$ represents the logistic regression, which is appropriate for estimating the probability as an outcome. Following the standard practice in RD design (Imbens and Lemieux (2008), Angrist and Pischke (2008), and Lee and Lemieux (2010)), we also set a bandwidth h, such that we estimate the treatment effects over $\Big|\text{Order\_Frac\_Hour}_i - 11\Big| \leq h$. For all of our analyses, we set the bandwidth to be 1 hour (from 10 am to 12 pm), which gives us a total of 16,710 orders to analyze. We note that our findings remain the same with different bandwidths, larger or smaller.

Including the τ_i and the interaction term directly yields the estimate of the treatment effect $\gamma =$

$$\left(\underset{\text{Order\_Frac\_Hour}\downarrow 11}{lim} E[Y(1)|\text{AFTER\_11} = 1] - \underset{\text{Order\_Frac\_Hour}\uparrow 11}{lim} E[Y(0)|\text{AFTER\_11} = 0] \right)$$

$$\text{or } \left(\underset{\tau\downarrow 0}{lim} E[Y(1)|\text{AFTER\_11} = 1] - \underset{\tau\uparrow 0}{lim} E[Y(0)|\text{AFTER\_11} = 0] \right)$$

where $Y(1)$ and $Y(0)$ are the outcome variables of orders placed after 11 am and before 11 am, respectively.

(iii) Duration of Leg t, (iv) Ratio for Leg t and (v) Average Action Time:

To estimate the impact of the 11 am same-day/next-day threshold on the duration of each leg of the delivery process, we have:

$$\text{Duration}_{i,t} = \alpha_t + \beta_t \tau_i + \gamma_t \text{AFTER\_11}_i + \delta_t \tau_i \cdot \text{AFTER\_11}_i + \zeta_t X_i + \epsilon_{i,t}$$

where $\text{Duration}_{i,t}$ represent the duration of the t-th leg ($t = 1, 2, 3$) in order i's delivery process, respectively.

The model for $\text{Ratio}_{i,t}$ is defined in the same way as the model for $\text{Duration}_{i,t}$ and the model for $\text{Average\_Action\_Time}_i$ is defined in the same way as the model for $\text{Duration}_{i,t}$, but for the whole delivery process, not for each leg.

4.4. Validation of assumptions

The RD framework requires two major assumptions. The assumptions are that (a) variables other than the assignment variable (in our case, AFTER_11) are continuous around the threshold, and (b) the individuals in the RD design cannot precisely manipulate the assignment variable. In this section we show that these assumptions hold in our setting.

(a) Covariate balance

To verify the first assumption, we check whether other covariates $X_{i,t}$ representing order and demographic characteristics change significantly at the discontinuity (i.e., at the 11 am threshold).

$$\lim_{\tau_i \downarrow 0} E[X_i(1)|\tau_i] \approx \lim_{\tau_i \uparrow 0} E[X_i(0)|\tau_i]$$

In Table 2.1 we show that the covariates are balanced between orders placed before 11 am (same-day) and those placed after 11 am (next-day). Numbers in the first row indicate the coefficients of the AFTER 11 variable, and the values directly beneath indicate the p-values of the coeffi-cients. Overall, most of the coefficients were insignificant, meaning there was no significant differ-ence in those characteristics between the same-day and next-day orders. However, the differences in Contains Gift variable ($p < 0.1$) and the Contains Quantity Discount variable ($p < 0.01$) were significant, although there is no clear evidence of what might be driving those factors, and the magnitudes of these coefficients are small.

Table 2.1 Covariate balance for: (1) Total Order Cost, (2) Number of SKUs, (3) Contains Gift, (4) Number of Items, (5) Contains Bundle Discount, (6) Contains Direct Discount, (7) Contains Quantity Discount, (8) Used a Coupon.

	List of covariates:							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AFTER_11	-7.0	0.005	0.2^{*}	0.01	-0.2	-0.1	0.2^{***}	0.1
	$p = 0.3$	$p = 0.8$	$p = 0.1$	$p = 0.9$	$p = 0.5$	$p = 0.2$	$p = 0.004$	$p = 0.5$
τ	11.0	0.01	-0.3^{*}	0.1	0.5^{*}	-0.1	-0.1	-0.3^{***}
	p = 0.2	p = 0.6	p = 0.1	p = 0.2	p = 0.1	p = 0.5	p = 0.2	p = 0.001
$\tau \cdot$ AFTER_11	-12.3	-0.01	0.2	-0.2^{**}	-0.3	0.3^{**}	-0.1	0.3^{**}
	$p = 0.3$	$p = 0.7$	$p = 0.4$	$p = 0.04$	$p = 0.4$	$p = 0.04$	$p = 0.4$	$p = 0.05$
Constant	125.4^{***}	1.1^{***}	-2.7^{***}	1.5^{***}	-3.6^{***}	1.1^{***}	-0.9^{***}	-1.1^{***}
	$p = 0.0$	$p = 0.0$	$p = 0.0$	$p = 0.0$	$p = 0.0$	$p = 0.0$	$p = 0.0$	$p = 0.0$

Note: *p<0.1; **p<0.05; ***p<0.01

(b) Assignment variable manipulation

To verify the second assumption, we check whether the outcome variable in the RD design is directly and deliberately manipulated by the agents in the system. The reason being that we want to ensure it is as if every observation around the cutoff was randomly assigned to either just above or just below the cutoff, like a randomized controlled trial (Lee and Lemieux (2010)). Thus, we don't want the agents to impose deliberate influence over the outcome variable.

Figure 2.3 shows the histogram of the hours of when the orders were placed, with the dashed line being a vertical axis equivalent to 11 am. Indeed, confirming our intuition, we see that there is a discontinuity in the volume of orders. There is an increase in customer demand leading right up to 11 am, with more customers preferring to have their orders delivered during the same day rather than wait for the next day. Despite the discontinuity in order volume, there is little evidence to believe that customers are placing the orders to directly and deliberately manipulate the delivery process. Since our outcome variables are focused on the delivery process itself, and not customer satisfaction, we don't necessarily care that customers can manipulate whether they place their orders right before or after the 11 am cutoff. Rather, we care that delivery and fulfillment center workers are not able to manipulate their workload around the 11 am cutoff. Given that the workers in charge of an order's delivery have no control over when orders are placed, we believe that this discontinuity will not violate the second assumption.

4.5. Results

In this section we present the results of our models for each outcome variable.

(i) Delayed and (ii) Late in Leg t

We first explore the relationship between time allotted and delay probability. To do this, we examine the plot of order time hour vs. $P\,(\text{Delay})$ in Figure 2.4. We see that there is an upward spike for $P\,(\text{Delay})$ occurring right before the 11 am cutoff, with the probability of an order being delayed falling dramatically with the increase in allotted time after 11 am. The probability of an order being delayed then slowly increases after the cutoff, finally reaching another peak right before 11 pm. We also examine how the probability of each leg being late is affected by allotted time. We see very similar results for all three legs, with $P\,(\text{Late leg}_t)$ spiking just before the 11 am cutoff and then falling with the increase in allotted time after the cutoff.

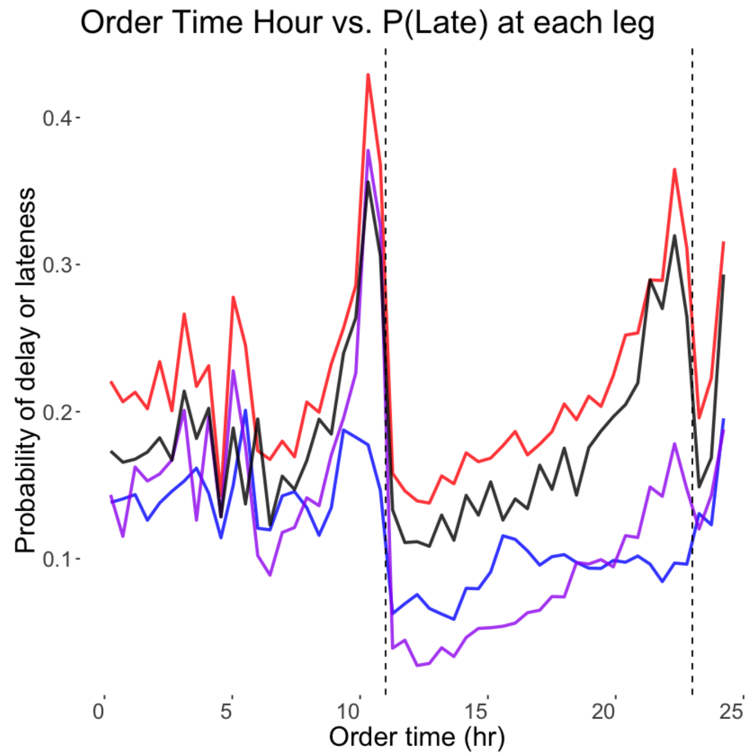

Figure 2.4 **Order Time vs. Probability of Delay and Lateness. The red line indicates $P(\text{Delay})$, i.e., the probability of the order not being delivered by the promised time. The black line indicates $P(\text{Late\_leg})$ for the first leg, purple line indicates $P(\text{Late\_leg})$ for the second leg, and the blue line indicates $P(\text{Late\_leg})$ for the third leg.**

Indeed, our regression discontinuity results from Table 2.2 confirm our visual observations. For P (Delay) and P (Late leg_t) occurring at leg t, we see that the log-odds of the AFTER 11_variable are all significant and negative. The coefficient on AFTER 11_where the outcome is the probability an order is delayed is -0.883 ($p < 0.01$) indicating that an increase in allotted time to deliver an order (due to switching from same to next-day delivery) significantly decreases the likelihood that the order is delayed. Similarly, the coefficients on AFTER 11 where the outcome variables are the probability of an order being late in the first leg, late in the second leg, and late in the third leg are -0.42, -2.38, and -1.00 (with all $p < 0.01$). This indicates that an increase in allotted time also decreases the probability that each individual leg of the order's delivery process is late.

(iii) Duration of Leg t

Table 2.2 RD model for probability of delay by leg for all orders placed between 10 am and 12 pm, with promise of same- and next-day delivery.

	Dependent variable:			
	is_delayed	late_first_leg	late_second_leg	late_third_leg
	(1)	(2)	(3)	(4)
AFTER_11	-0.88^{***}	-0.42^{***}	-2.38^{***}	-1.00^{***}
	(0.08)	(0.08)	(0.13)	(0.12)
τ	-0.64^{***}	-1.44^{***}	-0.54^{***}	-0.39^{***}
	(0.08)	(0.08)	(0.08)	(0.10)
AFTER_11 $\cdot\, \tau$	0.35^{**}	0.84^{***}	0.67^{***}	0.67^{***}
	(0.14)	(0.15)	(0.23)	(0.20)
Constant	-0.68^{***}	-1.23^{***}	-0.83^{***}	-1.81^{***}
	(0.04)	(0.04)	(0.04)	(0.05)
Observations	16,710	16,710	16,710	16,710
Log Likelihood	−9,441.53	−8,721.08	−7,372.47	−5,957.27
Akaike Inf. Crit.	18,891.05	17,450.17	14,752.95	11,922.54
Note:			*p<0.1; **p<0.05; ***p<0.01	

Next, we investigate how the time allotment affects the duration of each of the three individual legs in the delivery process. Our findings suggest that the effect of increasing allotted time is different for orders that ultimately arrive on-time and orders that end up delayed. These findings are visually evident in the plots in Figure 2.5, which show the relationship between order hour of the day and duration of each leg for on-time and delayed orders.

For the on-time orders, from Table 2.3 we see that the next-day orders (i.e., orders placed post-11 am) on average took AFTER $11_1 = 3.46$ ($p < 0.01$), AFTER $11_2 = 10.85$ ($p < 0.01$), AFTER $11_3 = 3.49$ ($p < 0.01$) hours longer than the same-day orders for the 1st, 2nd, and 3rd legs, respectively. On the other hand, we see a slightly different effect for delayed orders (columns 4 to 6). For the first leg (i.e., $\mathrm{D}_1^{\mathrm{Delay}}$), we see that the first leg duration for the delayed next-day orders is substantially larger AFTER $11_1 = 20.89$ ($p < 0.01$) than the delayed orders that are promised same-day delivery. However, we see that the difference between same-day and next-day is insignificant for the second leg AFTER $11_2 = -0.72$ ($p = 0.169$) of the delayed orders. For the third leg of the delayed orders, we see that the effect is similar to that of the on-time orders ($\text{AFTER\_11}_3 = 3.60$ ($p < 0.01$)).

For the regression results with full list of other control variables X_i, please see Table A2.5. We note that, for all the regression discontinuity analyses and the outcome variables in this section, our findings are substantively similar even when other control variables are included.

Table 2.3 RD model for duration of each leg for all orders placed between 10 am and 12 pm, with promise of same- and next-day delivery.

	Dependent variable:					
	D_1^{Ontime}	D_2^{Ontime}	D_3^{Ontime}	D_1^{Delay}	D_2^{Delay}	D_3^{Delay}
	(1)	(2)	(3)	(4)	(5)	(6)
AFTER_11	3.46^{***}	10.85^{***}	3.49^{***}	20.89^{***}	-0.72	3.60^{***}
	(0.05)	(0.14)	(0.14)	(1.60)	(0.52)	(0.94)
τ	-1.00^{***}	-0.02	-0.01	-4.14^{***}	0.79^{*}	-0.55
	(0.07)	(0.18)	(0.18)	(1.34)	(0.43)	(0.78)
$\text{AFTER\_11} \cdot \tau$	-0.18^{*}	-0.39	0.52^{**}	-8.59^{***}	-0.23	5.42^{***}
	(0.10)	(0.25)	(0.25)	(2.81)	(0.91)	(1.65)
Constant	0.005	3.73^{***}	2.39^{***}	7.54^{***}	14.82^{***}	8.11^{***}
	(0.04)	(0.09)	(0.09)	(0.75)	(0.24)	(0.44)
Observations	11,786	11,786	11,786	4,924	4,924	4,924
R^2	0.43	0.68	0.21	0.05	0.001	0.03
Adjusted R^2	0.43	0.68	0.21	0.05	0.0002	0.03
Residual Std. Error	1.44	3.63	3.65	23.44	7.62	13.76
F Statistic	$2{,}909.07^{***}$	$8{,}394.46^{***}$	$1{,}035.07^{***}$	95.02^{***}	1.29	54.88^{***}
Note:						*p<0.1; **p<0.05; ***p<0.01

(iv) Ratio for Leg t

Because there may be inherent differences among the three legs in their duration and how they process their orders, we performed an RD analysis with the ratio (duration/time allotted) as an outcome variable to normalize duration of individual legs. In our analysis, having a positive (negative) coefficient would indicate that next-day orders – compared to the same-day orders – are taking a disproportionately longer (shorter) time when we take the increased amount of allotted time into account.

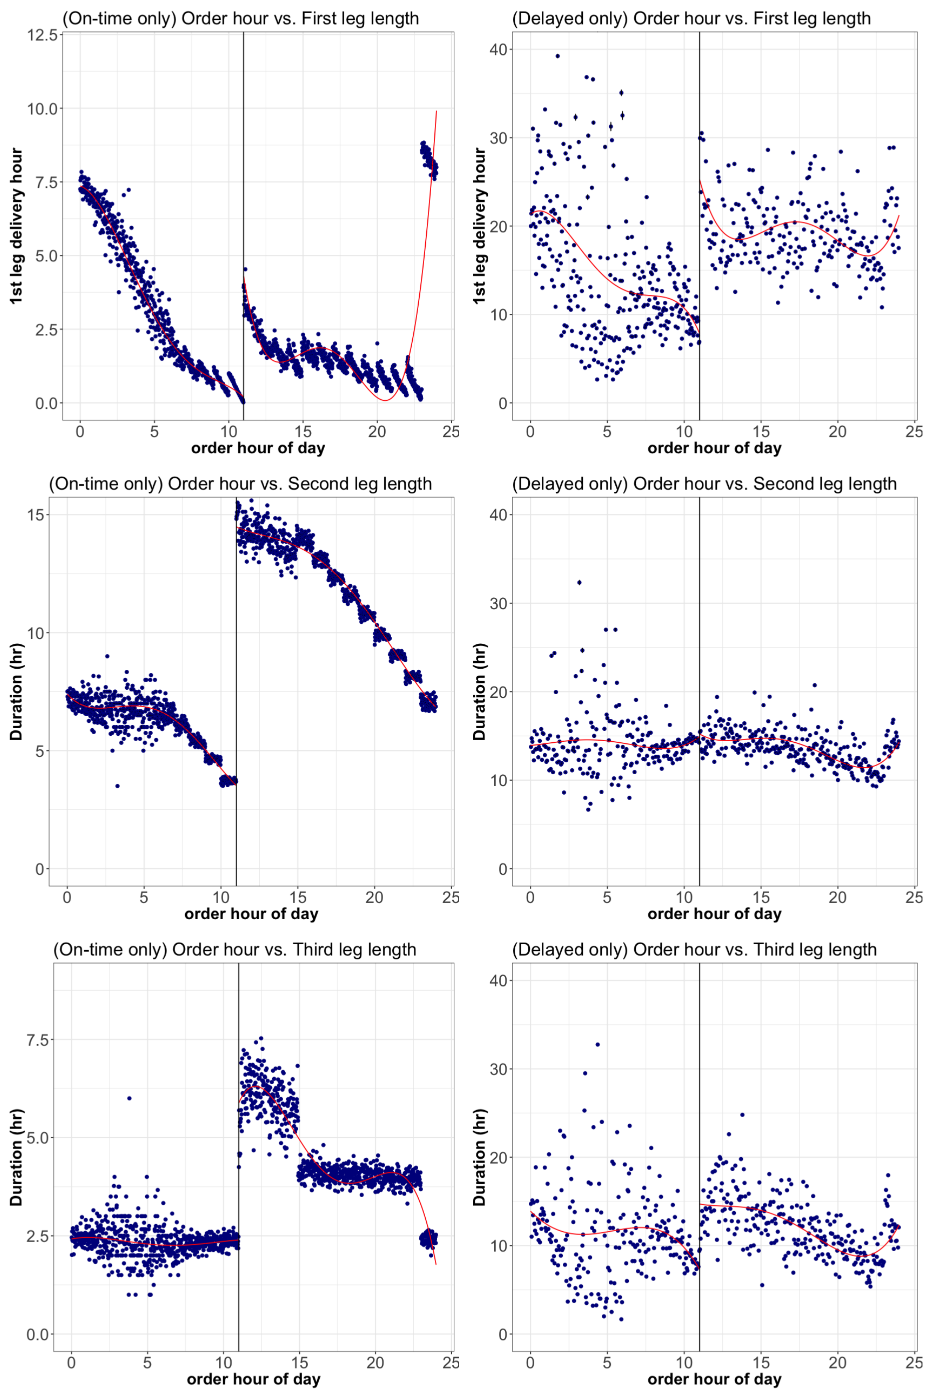

Leg 1

Leg 2

Leg 3

On-Time Orders **Delayed Orders**

Figure 2.5 Order Hour vs. Leg Duration for each of the three legs (row) and for on-time and delayed orders (columns). Vertical lines (i.e., $x = 11$) represent the threshold for same-day vs. next-day orders.

For on-time orders, the results in Table 2.4 suggest that the AFTER 11 coefficients for the three legs are all positive and significant (AFTER $11_1 = 0.12$ ($p < 0.01$), AFTER $11_2 = 0.23$ ($p <$

0.01), AFTER_11$_3 = 0.03$ ($p < 0.01$). This indicates that the next-day orders are taking disproportionately longer, even when leg durations are normalized with allotted time. For the delayed orders, the ratio of the first leg increases with the additional allotted time (AFTER_11$_1 = 0.43$ ($p < 0.01$)), but we see that the coefficients for the second and third legs are negative (AFTER_11$_2 = -0.64$ ($p < 0.01$), AFTER_11$_3 = -0.21$ ($p < 0.01$)), indicating the opposite effect.

Table 2.4 RD model for (duration to time allotment) ratio for all orders placed between 10 am and 12 pm, with promise of same- and next-day delivery.

	Dependent variable:					
	$RATIO_1^{\text{Ontime}}$ (1)	$RATIO_2^{\text{Ontime}}$ (2)	$RATIO_3^{\text{Ontime}}$ (3)	$RATIO_1^{\text{Delay}}$ (4)	$RATIO_2^{\text{Delay}}$ (5)	$RATIO_3^{\text{Delay}}$ (6)
AFTER_11	0.12***	0.23***	0.03***	0.43***	−0.64***	−0.21***
	(0.002)	(0.01)	(0.01)	(0.10)	(0.03)	(0.06)
τ	−0.07***	0.02***	0.01*	−0.25***	0.14***	0.004
	(0.003)	(0.01)	(0.01)	(0.09)	(0.03)	(0.05)
AFTER_11 $\cdot\, \tau$	0.03***	−0.01	0.01	−0.18	−0.10*	0.19*
	(0.004)	(0.01)	(0.01)	(0.18)	(0.06)	(0.10)
Constant	0.001	0.29***	0.18***	0.59***	1.14***	0.62***
	(0.001)	(0.003)	(0.004)	(0.05)	(0.02)	(0.03)
Observations	11,786	11,786	11,786	4,924	4,924	4,924
R^2	0.35	0.44	0.03	0.004	0.19	0.004
Adj. R^2	0.35	0.44	0.02	0.003	0.19	0.003
Resid. SE	0.05	0.14	0.14	1.50	0.48	0.82
F Statistic	2,116.78***	3,064.93***	100.91***	6.56***	393.06***	6.70***
Note:						*p<0.1; **p<0.05; ***p<0.01

(v) Average Action Time

Although the dataset we were given by JD.com does not contain customer satisfaction or ratings data, we use the findings in Bray (2019) to infer the effect of time allotment on customer satisfaction. This result is based on the research idea "peak-end effect" presented in Varey and Kahneman (1992), that the later moments of an activity rather than the earlier moments are most salient to the customers. In the context of e-commerce delivery, Bray (2019)'s finding suggests that a particularly slow (fast) delivery activity towards the end will leave a negative (positive) impression on the customers, thus they will be less (more) satisfied.

Within our data sample, the mean average action time across all orders is 0.46 with a standard deviation of 0.16. This is similar to what Bray (2019) found in a dataset of online deliveries by Cainiao Network, where the mean average action time was 0.49, with a standard deviation of 0.13, indicating that his results are likely applicable to the JD.com dataset.

We look at orders around the same-day next-day delivery cutoff to understand how the discontinuous increase in time allotment affects their average action time. We find that increasing the time allotted for an order results in a positive and significant increase in the order's average action time, with the coefficient of AFTER 11 $= 0.114$ ($p < 0.01$) (see Table 2.5 for full results). These results are robust across both delayed and on-time orders. Bray (2019) finds that delaying the average action time of an order (thus increasing its average action time) increases the order's delivery score, which is a metric of customer satisfaction around the delivery process. Therefore we find that, although the same-day next-day cutoff disproportionately increases the time it takes an order to be delivered, it alters the delivery process in such a way that the average action time is pushed further back. And if the relationship between delayed average action time and delivery scores holds true in our setting, it is plausible that these orders with increased allotted time will also have higher customer satisfaction, even though they take longer to be delivered.

5. Delay Prediction Model Using Random Forest

5.1. Approach

We consider the next set of naturally arising questions: (a) can we predict which orders will be delayed, given the order and user characteristics? (b) which features are the most important in predicting delays? and (c) can knowing whether an earlier leg of the delivery process is late increase the accuracy of delay predictions? To answer the questions above, we use the random forest algorithm for building regression and classification trees. Random forest is a bagging algorithm, which

Table 2.5 RD model for average action times for all orders placed between 10 am and 12 pm, with promise of same- and next-day delivery.

	Dependent variable:		
	Average action time		
	All orders	On-time orders	Delayed orders
AFTER_11	0.114^{***}	0.136^{***}	0.150^{***}
	(0.005)	(0.004)	(0.015)
τ	-0.095^{***}	-0.098^{***}	-0.051^{***}
	(0.006)	(0.005)	(0.012)
AFTER_11 $\cdot\, \tau$	0.041^{***}	0.060^{***}	-0.058^{**}
	(0.009)	(0.007)	(0.026)
Constant	0.364^{***}	0.316^{***}	0.453^{***}
	(0.003)	(0.003)	(0.007)
Observations	16,710	11,786	4,924
R^2	0.039	0.146	0.027
Adjusted R^2	0.039	0.145	0.027
Residual Std. Error	0.153 (df = 16706)	0.101 (df = 11782)	0.215 (df = 4920)
F Statistic	225.200^{***} (df = 3; 16706)	669.251^{***} (df = 3; 11782)	45.775^{***} (df = 3; 4920)
Note:			*p<0.1; **p<0.05; ***p<0.01

iteratively explores different subsets of predictors available (James et al. (2013)). Such an iteration process is designed to identify the most important predictors, while simultaneously dropping correlated but less important predictors (which is important for problems arising in analyzing retail data, many of whose features are correlated), thus reducing the variance and increasing the reliability of the resulting trees.

Formally, as described in (Breiman 2001), for each order n we have labels $y_n, \forall n = 1, 2, ..., N_{orders}$ and features $\vec{x_n}$, together comprising the dataset $\left\{(y_1, \vec{x_1}), (y_2, \vec{x_2}), ..., (y_{N_{orders}}, \vec{x}_{N\,orders})\right\}$. For each iteration $k, \forall k = 1, 2, ..., K$ of random forest, we select a subset of features $\Theta_k \subset \Theta$ to construct classification trees $\mathcal{T}_1(\vec{x}|\Theta_1), \mathcal{T}_2(\vec{x}|\Theta_2), ..., \mathcal{T}_k(\vec{x}|\Theta_k), ...,$ with classifiers $h_1(\vec{x}|\Theta_1), h_2(\vec{x}|\Theta_2), ..., h_k(\vec{x}|\Theta_k), ...,$, which give classification outcomes $\hat{y_{k,1}}, \hat{y_{k,2}}, ..., \hat{y_{k,n}}, ..., \hat{y}_{k,N_{orders}}$ for order n and k-th iteration. Subsequently, the average marginal performance of the random forest model, as described in (Breiman 2001) is:

$$\hat{m}(\vec{x}, y) \equiv \left[\hat{P}_k(h_k(\vec{x}|\Theta) = y) - \max_{j \neq y}\big(\hat{P}_k(h_k(\vec{x}|\Theta) = j)\big)\right]$$

where the above measures the difference between the number of votes for the "most popular" class (given by the k iterations) and the number of votes for the remaining classes. and the generalization error ϵ is defined as:

$$\epsilon = P(\hat{m}(\vec{x}, y) < 0)$$

5.2. Model and Feature Selection

We ran a random forest model through R's "randomForest" package. For the parameters, we run the algorithm with the $50-50$ train-test split. Furthermore, we set the number of trees and max nodes as 200.

The outcome we are trying to predict is whether an order will be delayed and ultimately arrive at the customer's house after its promised delivery time. This outcome is a binary variable that takes value 1 if an order is delayed and 0 if it arrives on time. We ran and analyzed the random forest models with three different sets of features. The first model, our baseline model, includes a set of features containing order and user characteristics. These features are all determined when the order is placed, prior to the beginning of the delivery process. Some of these features, such as customer's gender and city level (e.g. more rural or urban), are fixed variables. Other features vary with the order and depend on customers' decisions, such as the time of day the order is placed and the total cost of the order. We also include features that are determined by JD.com following the placement of the order, such as which warehouse will fulfill the order. See Table A2.6 in the Online Appendix for the full list of features.

The second and the third models use all the features mentioned previously in the first model, as well as partial information about the earlier legs of the delivery process. The main idea here is to study whether adding information for in-progress deliveries on the earlier legs in the delivery process can significantly increase the model's accuracy in predicting delays. Specifically, the second model contains all the features in the first model plus the features related to the first leg, i.e., how long the first leg took, what time of day the first leg was completed, and how long the order has left to be delivered after the first leg. Finally, the third model contains all the features in the baseline model, as well as additional information about the first and the second leg (i.e., the length of the second leg, the time of day the second leg was completed, and how much time is remaining after the second leg).

5.3. Results

Table 2.6 contains the performance results of the above-mentioned three random forest models. Our 'No Information Rate' was 77.5%, which was the base-rate percentage of on-time orders in the test set. As evidenced by 'P-value [Acc > NIR]' row (where NIR is determined by Naïve Bayes Classifier with the error rate $\frac{1}{n^2}\sum_{i=1}^{n}\sum_{p=1}^{n}\mathcal{L}(y_i, f(x_p))$ for loss function $\mathcal{L}$, label y_i and features x_p), the baseline model performed better than the 'No Information Rate' (Accuracy CI of Baseline: $(0.797, 0.803)$). However, the model also containing the first leg information (duration, ship out time hour, and time left after the first leg) performed significantly better than the baseline prediction model (Accuracy CI with 1st leg information: $(0.915, 0.919)$). Adding both the first and the second leg information to the model also significantly increased prediction accuracy relative to the baseline model, but the performance difference between the second model (with 1st leg information) and the third model (with both 1st and 2nd leg) was nominal (Accuracy CI with 1st & 2nd leg information: $(0.936, 0.940)$).

Subsequently, Table 2.7 shows the list of top 15 features for each of the three models. The importance of the features are determined by the Gini impurity (Breiman et al. (1984), Raileanu and Stoffel (2004)), and we present the list of features in descending order of the Gini impurity measure in Table 2.7. Across the three models, we see that the delivery characteristics inherently related to the order time – such as order time day or time allotted - have greater feature importance than order characteristics (e.g., whether an order contains a gift or a coupon) or user characteristics

Table 2.6 Random forest model prediction performance results for predicting delay. The results in the first column are for the baseline model (without any duration information), and those in the second and the third column are for the model with first leg information, and first & second leg information, respectively. The top cells, i.e., (Prediction) vs. (True Values), show the confusion matrix (0 = not delayed, 1 = delayed), the second cells show overall accuracy of the model, as well as 95% CI.

Confusion Matrix and Statistics (Predicting Delays)						
	Baseline		1st leg info.		1st + 2nd leg info.	
	(True Values)		(True Values)		(True Values)	
(Prediction)	0	1	0	1	0	1
0	48926	12397	49090	5105	49064	3741
1	314	1890	150	9182	176	10546
Accuracy	**0.800**		**0.917**		**0.938**	
95% CI	**(0.797, 0.803)**		**(0.915, 0.919)**		**(0.936, 0.940)**	
No Information Rate (NIR)	0.775		0.775		0.775	
P-Value [Acc>NIR]	<2.2E-16		<2.2E-16		<2.2E-16	
Kappa	0.180		0.729		0.806	
Mcnemar's Test P-Value	<2.2E-16		<2.2E-16		<2.2E-16	
Sensitivity	0.994		0.997		0.996	
Specificity	0.132		0.643		0.738	
Pos Pred Value	0.798		0.906		0.929	
Neg Pred Value	0.858		0.984		0.984	
Prevalence	0.775		0.775		0.775	
Detection Rate	0.770		0.773		0.772	
Detection Prevalence	0.965		0.853		0.831	
Balanced Accuracy	0.563		0.820		0.867	

(e.g., user's marriage status, age, education). As expected, in the second and the third model, the information about the earlier legs (i.e., first and second leg) show substantially higher variable importance than other features. The full list of features by predictive importance is available in Table A2.6 in the Online Appendix. This finding suggests that including information about the earlier legs in the prediction models – whether real-time or ex-post - can substantially improve the accuracy and the overall prediction performance. We note that other methods such as logistic

Table 2.7. Top 15 Random forest model prediction performance results for predicting delay.

Full List of Top Features (and Corresponding Gini Impurity Values)					
Baseline		1st leg info.		1st + 2nd leg info.	
order_time_day	872.68	time_left_after_first_leg	6995.82	time_left_after_second_leg	6555.10
time_allotted	546.66	first_leg_length	3513.60	time_left_after_first_leg	4605.76
dc_ori	293.06	ship_out_time_frac_hour	589.04	first_leg_length	2056.99
order_time_frac_hour	288.14	time_allotted	391.39	second_leg_length	847.88
dc_des	264.12	order_time_day	381.30	order_time_day	198.73
order_time_day_of_week	215.16	order_time_frac_hour	295.42	time_allotted	191.73
total_order_cost	212.23	dc_ori	95.78	ship_out_time_frac_hour	179.05
city_level	77.62	after_11	95.12	arr_station_time	160.20
type	68.71	dc_des	92.34	order_time_frac_hour	140.38
after_11	48.77	total_order_cost	75.02	dc_ori	57.00
user_level	40.50	order_time_day_of_week	58.12	dc_des	51.71
contains_quantity_discount	39.27	city_level	33.74	total_order_cost	39.89
num_gifts	32.02	age	21.14	city_level	28.26
age	31.37	type	20.98	order_time_day_of_week	24.45
contains_gift	31.00	user_level	18.18	after_11	20.26

regression and gradient-boosted trees (XGBoost) were also used, and the results were substantively similar.

6. Building and Incorporating Load-Related Features

Finally, we augment the above random forest model by adding additional load-related features. Delivery information in this JD.com dataset could be formulated as a fluid model in queueing as described in (Hall 1991, Newell 2013). In this example, we focus on the origin warehouse (denoted as ORI), where the package completes its first leg's journey (i.e., ship out to the $DEST$ warehouse) and begins its second leg.

Let $N_i(t)$ denote the number of packages at the origin warehouse at hour t (from hour $t-1$ to t), where $i=1,...,N_{ORI}$ and $t=1,2...,T$. We assume that t is a rolling time period, so any unprocessed packages from 11:59 pm the previous day will need to be processed the next day. Further, let $A_i(t)$ denote the number of packages arriving at the origin warehouse i, and let $B_{i,j}(t)$ denote the number of packages leaving origin warehouse i to destination warehouse j. For each hour t (from hour $t-1$ to t) and origin warehouse i, we can formulate the following recursive framework:

$$N_i(t) = A_i(0) \text{ for } t=0$$

$$N_i(t) = \Big(N_i(t-1) + A_i(t) - \sum_{j \in DEST} B_{i,j}(t)\Big)^+ \text{ for } \forall t=1,...,T-1 \text{ and } \forall i=1,...,N_{ORI}$$

$$N_i(t) = \Big(N_i(t-1) - \sum_{j \in DEST} B_{i,j}(t)\Big)^+ \text{ for } t=T \text{ and } \forall i=1,...,N_{ORI}$$

With the fluid model framework described in (Hall 1991, Newell 2013), we can define the arrival and processing rates (denoted $\lambda_i(t)$ and $\delta_{i,j}(t)$, respectively) by differentiating $A_i(t)$ and $B_{i,j}(t)$ with respect to t:

$$\lambda_i(t) = \dot{A_i}(t) = \frac{dA_i(t)}{dt}$$

$$\delta_{i,j}(t) = \dot{B_{i,j}}(t) = \frac{dB_{i,j}(t)}{dt}$$

where the difference $\Big(\lambda_i(t) - \sum_{j \in DEST} \delta_{i,j}(t)\Big)$ is equivalent to the net package flow rate, $\dot{N_i}(t)$.

These ordinary differential equations, then, can be solved with the following approximation formula for time period between $[t-\Delta t, t]$ where $t \geq \Delta t$:

$$\begin{aligned}\sum_{\tau=t-\Delta t}^{t} N_i(\tau) &= N_i(t-\Delta t) + \Big(\sum_{\tau=t-\Delta t}^{t} \frac{\lambda_i(\tau)}{\Delta t}\Big)\Delta t - \Big(\sum_{\tau=t-\Delta t}^{t} \frac{\delta_i(\tau)}{\Delta t}\Big)\Delta t \\ &= N_i(t-\Delta t) + \sum_{\tau=t-\Delta t}^{t} \lambda_i(\tau) - \sum_{\tau=t-\Delta t}^{t} \delta_i(\tau)\end{aligned}$$

Here, $\sum_{\tau=t-\Delta t}^{t} N_i(\tau)$ represents the number of packages in time period between $[t-\Delta t, t]$. $\sum_{\tau=0}^{t-\Delta t} N_i(\tau)$ represents the number of packages between $[t-\Delta t, t]$. The third and fourth terms $\Big(\sum_{\tau=t-\Delta t}^{t} \frac{\lambda_i(\tau)}{\Delta t}\Big)$ and $\Big(\sum_{\tau=t-\Delta t}^{t} \frac{\delta_i(\tau)}{\Delta t}\Big)$ represent the average arrival rate and average departure rate between times $[t-\Delta t, t]$, respectively.

Furthermore, notice that the number of departures depend on how many packages are currently in the system. Let $\mu_{i,j}(t)$ be the package processing rate at hour t done by one worker from warehouse i to j, and $s_{i,j}(t)$ be the number of available workers in charge of sending packages from warehouse i to j. Then, $\Big(\mu_{i,j}(t) \cdot s_{i,j}(t)\Big)$ represents the package processing capacity, i.e., the maximum number of packages that can be processed in hour t. If there are more packages than the max capacity, i.e., $\Big(N_i(t-1)+\lambda_i(t)\Big) \geq \sum_j \Big(\mu_{i,j}(t) \cdot s_{i,j}(t)\Big)$, then the number of packages processed for warehouse i at hour t will be $\sum_j \Big(\mu_{i,j}(t) \cdot s_{i,j}(t)\Big)$. Otherwise, the number of processed packages will be $\Big(N_i(t-1)+\lambda_i(t)\Big)$, and there will be zero packages left over until the next hour.

Thus, $\delta_i(\tau)$ can be expressed as:

$$\sum_{\tau=t-\Delta t}^{t} \delta_i(\tau) = \sum_{\tau=t-\Delta t}^{t} min\Big[\sum_{j \in DEST} \Big(\mu_{i,j}(\tau) \cdot s_{i,j}(\tau)\Big), \Big(N_i(\tau-1)+\lambda_i(\tau)\Big)\Big]$$

As a result, the number of packages between the hours $t-\Delta t$ and t becomes:

$$\sum_{\tau=t-\Delta t}^{t} N_i(\tau) = N_i(t-\Delta t) + \sum_{\tau=t-\Delta t}^{t} \lambda_i(\tau) - \sum_{\tau=t-\Delta t}^{t} min\Big[\sum_{j \in DEST} \Big(\mu_{i,j}(\tau) \cdot s_{i,j}(\tau)\Big), \Big(N_i(\tau-1)+\lambda_i(\tau)\Big)\Big]$$

When substituting arrival and departure rates $\lambda_i(t)$ and $\sum_{j\in DEST}\mu_{i,j}(t)$ with the inverse of the inter-arrival and inter-departure times ($\theta_{\lambda_i}^{-1}(t)$ and $\theta_{\mu_{i,j}}^{-1}(t)$) in the dataset, we have:

$$\sum_{\tau=t-\Delta t}^{t} N_i(\tau) = N_i(t-\Delta t) + \sum_{\tau=t-\Delta t}^{t} \theta_{\lambda_i}{}^{-1}(\tau) - \sum_{\tau=t-\Delta t}^{t} min\Big[\sum_{j\in DEST}\Big(\theta_{\mu_{i,j}}{}^{-1}(\tau)\cdot s_{i,j}(\tau)\Big),\Big(N_i(\tau-1)+\theta_{\lambda_i}{}^{-1}(\tau)\Big)\Big]$$

6.1. Random forest with additional load features

We run a random forest model similar to the earlier ones run in Section 5.3. The additional load-related features are: (i) number of packages at time t at warehouse i, $N_i(t)$ as recursively calculated in the above equation, (ii) mean and variance of inter-arrival times at time t at warehouse i, $\theta_{\lambda_i}^{-1}(t)$, (iii) mean and variance of inter-departure times, $\theta_{\mu_{i,j}}^{-1}(t)$, and (iv) the dispersion measure, defined as $\mathcal{I}(t) = Var(x(t))/\overline{x(t)}$ for variable $x(t)$, for inter-arrival and inter-departure times. We did not have visibility into the number of servers at each hour, so we assume $s_{i,j}$ to be 1 for this dataset. The outcome again is a binary indicator for whether a shipment is delayed, which takes value 1 if an order is delayed. Similarly as the random forest models in Section 5.3, we ran the random forest algorithm with the 50 - 50 train-test split, with the number of trees and max nodes set as 200.

Table 2.8. compares the results of the random forest models with three different sets of features: baseline (order and user information), baseline & 1st leg information, and baseline & load-related features. Naïve Classifier results still remain at 77.5%; the accuracy level for the baseline, baseline & 1st leg, and baseline & load-related features were 80.0% (95% CI: $[79.7\%, 80.3\%]$), 91.7% (95% CI: $[91.5\%, 91.9\%]$), and 84.9% (95% CI: $[84.6\%, 85.2\%]$), respectively. Results suggest that including load-related features achieves an accuracy level significantly higher than that of the baseline model but is lower than the model with 1st leg information.

There may be couple of reasons why including load-related features do not perform as well as the model with 1st leg information. First, many delays may not be directly attributable to overload. Given there is no drastic/abrupt drop in service levels, overload will render the duration to increase incrementally, proportionally, and often linearly. Figure 2.4, however, shows a sharp, non-overlapping difference in duration between on-time orders and delayed orders. For the first-leg duration, for instance, on-time orders take about 4 hours to complete and no longer than 8 hours, whereas delayed orders average about 20 hours to complete (especially when the duration jumps around 11 am). This indicates that the delay may not be necessarily attributed to overload alone – and may instead be attributed to other factors, such as item stock-out. Second, lower accuracy with load-related features may be attributed to the lack of service level information (i.e., number of servers/labor units available at each hour). Without the service level information, it is difficult to predict how variable service levels are throughout the day. Lastly, there may be other ways to improve the model, e.g., through constructing an ensemble model that combines a number of other ML methods that give a better fit.

Table 2.8 Random forest model prediction performance results for predicting delay. The results in the first column are for the baseline model (without any duration or load information), those in the second column are for models with first leg information, and those in the third column are with warehouse package load information. The top cells, i.e., (Prediction) vs. (True Values), show the confusion matrix (0 = not delayed, 1 = delayed), the second cells show overall accuracy of the model, as well as 95% CI. Overall, the results show that accuracy level with the load information is significantly higher than the performance of the baseline model, but the accuracy is lower than the one with 1st leg information.

Confusion Matrix and Statistics (Predicting Delays)						
	Baseline		1st leg info.		load info.	
	(True Values)		(True Values)		(True Values)	
(Prediction)	0	1	0	1	0	1
0	48926	12397	49090	5105	45623	8471
1	314	1890	150	9182	488	4781
Accuracy	**0.800**		**0.917**		**0.849**	
95% CI	**(0.797, 0.803)**		**(0.915, 0.919)**		**(0.846, 0.852)**	
No Information Rate (NIR)	0.775		0.775		0.775	
P-Value [Acc>NIR]	<2.2E-16		<2.2E-16		<2.2E-16	
Kappa	0.180		0.729		0.446	
Mcnemar's Test P-Value	<2.2E-16		<2.2E-16		<2.2E-16	
Sensitivity	0.994		0.997		0.989	
Specificity	0.132		0.643		0.361	
Pos Pred Value	0.798		0.906		0.843	
Neg Pred Value	0.858		0.984		0.907	
Prevalence	0.775		0.775		0.777	
Detection Rate	0.770		0.773		0.769	
Detection Prevalence	0.965		0.853		0.911	
Balanced Accuracy	0.563		0.820		0.675	

7. Managerial Implications, Limitations, and Conclusions

In summary, we have explored the relationship between time allotment and delivery characteristics – namely delay, lateness, duration, ratio, and action time – using regression discontinuity design and machine learning models. The regression discontinuity analyses revealed that the same-day orders have shorter delivery durations than next-day orders, even when the durations are normalized by time allotted. The same-day orders are also shown to have lower action times, implying that their delivery takes relatively longer for the later stages (e.g., the third leg) compared to next-day orders. Furthermore, the same-day orders are shown to have higher probabilities of being delayed or late compared to the next-day orders. Finally, we fitted different versions of random forest models to predict delays; we find that partial information about earlier legs (e.g., first leg's duration and the amount of time remaining after the first leg) is valuable in predicting delays, as it substantially boosts the performance of our machine learning models. In addition, we develop and incorporate load-related features using fluid model in queueing, as described in (Hall 1991, Newell 2013). The results suggest that incorporating load-related features also can boost the model performance – but not to the same extent as incorporating partial information about earlier legs.

7.1. Managerial takeaways

First, our findings from the regression discontinuity results suggest that there are upsides (lower probabilities of delay and lateness, as well as lower action times) and downsides (disproportionately longer time duration) of increasing allotted time. When we look at the histogram of order volume by hour depicted in Figure 2.3, we see that 10-11 am is when the volume of orders is the highest. At the same time, however, 10-11 am is also the hour with the least time allotted (14 hours) for the company to ship out from the warehouse, truck the packages, and deliver the packages to the customers. Indeed, although these orders placed at 10-11 am takes less time to deliver, they are shown to have an approximately 40% probability of delay, suggesting that these orders might be rushed to get to the customers. Figure 2.5 illustrates similar patterns for orders placed between 10-11 pm, when the orders are placed right before the nighttime cutoff (recall that orders placed between

11 am and 11 pm are promised to be delivered by 3 pm next day). These orders are associated with shorter duration (both raw and normalized), but higher probabilities of delays associated with decreased time allotment. Finally, the results suggest that orders with less time allotted have lower average action times, which may indicate lower customer satisfaction (according to findings from Bray (2019)). For these time-constrained orders (i.e., orders placed at 10-11 am or 10-11 pm), customer satisfaction may be even lower due to higher probabilities of delays and lateness.

Obviously, the value of the tradeoff between the above-mentioned upsides and downsides would depend on the company's overall objective. It is well-established in operations literature (e.g., Kim and Yano (1994), Pinedo (2012)) that minimizing the number of delayed orders is a different objective than minimizing the average length of delay, for instance. In any case, an online retailer seeking to improve their delivery outcomes could employ several strategies, such as (a) adjusting staffing and trucking requirements to account for increased demand, (b) rather than constraining the order to be delivered on the same day (which could be as little as 14 hours), have the orders delivered in 24 hours so that the orders can still be labeled as "one-day delivery", except there would be more time available to fulfill the orders, and (c) prioritizing shipping out orders that are likely to be delayed due to lower time allotments, rather than simply following a FIFO process. On the other hand, behaviorally, following a standardized process (such as FIFO) is beneficial in reducing the cognitive load of the workers (Ibanez et al. 2018). Thus, retailers should ensure not to introduce an order assignment policy that increases workers' mental overhead.

Second, our machine learning models have illustrated the value of including information about the earlier legs in substantially improving the delay prediction model. Although the random forest models presented in the paper could be improved with other prediction algorithms, we have demonstrated that machine learning is valuable for accurately detecting early any orders at risk of being delayed.

To further substantiate the value of our delay prediction models, we quantify the concept of "recoverability". We deem an order to be recoverable after the first leg if it can still theoretically be delivered on time to the customer, even after having a late first leg. Recoverable is a binary variable equal to 1, if the time remaining for the order to be delivered after the first leg (Promise Time - Ship Out Time) is greater than the 1st percentile (fastest) of Second Leg Duration + Third Leg Duration of the orders placed at the same hour of the day. In other words, if an order has taken a long time to complete the first leg, it is recoverable if it still has enough time to complete the two remaining legs. If managers can accurately identify orders that are likely to be delayed but are also recoverable, they can then divert extra resources to ensuring these orders complete the rest of their delivery process on time.

We examine how our models perform in identifying recoverable delayed orders on a 50%-50% train-test split. The test dataset contains 14,287 delayed orders, of which 8,185 are recoverable. We find that our prediction model trained solely on order and customer characteristics known by the time the order is placed (baseline) predicts 2,204 orders to be delayed, of which 1,890 orders are correctly predicted as delayed, and 878 are also incidentally recoverable. In other words, if the managers in JD.com employ this particular model in the early detection of delayed orders, they would identify 10.73% ($\frac{878}{8185}$) of all delayed and recoverable orders. Subsequently, if JD.com increased resources for these orders throughout the delivery process, they could prevent them from arriving past the promised time, thus increasing customer satisfaction. Therefore, if JD.com applies a blanket policy to increase resources for any of the 2,204 orders predicted to be delayed, 39.84% ($\frac{878}{2204}$) of the orders JD.com tries to help could theoretically make it to the customer on time.

If JD.com were to consider using a prediction model that incorporates features from the first leg of the delivery process, they would be able to improve their accuracy in identifying delayed and recoverable orders. The model that adds features around the first leg of the delivery process was able to correctly predict 9,156 orders as being delayed, of which 3,054 are also recoverable. Therefore, this model would give JD.com the capability to identify around 37.31% ($\frac{3054}{8185}$) of all delayed and recoverable orders. Given that at this point the retailer has observed the length of the first leg, they will then be able to identify exactly which orders are recoverable and which are not (as recoverability is defined solely based on the length of the first leg and the allotted time). Thus, after the first leg has completed, JD.com can then use this targeted model to increase resources for all 8,226 orders that it predicts will be delayed and are recoverable. This policy might actually benefit 37.13% ($\frac{3054}{8226}$) of the orders which are given increased resources. Although this percentage of benefited orders is slightly lower than under the previous prediction model, the absolute number of orders that will be benefited (and might make it to the customer on time due to this intervention) has increased. However, given that predictions from this model will be made after the first leg has completed, there are fewer legs of the delivery remaining in which JD.com can intervene and divert resources, meaning that it might be more difficult to recover orders after the first leg has finished. An online retailer considering using a machine learning model to predict delayed and recoverable orders should balance the cost of diverting resources to an order that cannot be recovered with the benefits of getting an order to a customer on-time rather than after it was expected. The retailer should also consider at what point in the delivery process allows them the most freedom to intervene and increase resources for an order that might otherwise be delayed.

Lastly, managers could also incorporate load-related features in their prediction models to boost the model accuracy. Of course, how beneficial load-related features are depends on how often delays are attributable to overload in warehouses. In cases where warehouses do not experience significant delivery delays from package overload, load-related features may not contribute much to boosting the overall model performance.

7.2. Limitations and future directions

We take note of some limitations, which in turn could turn into opportunities for future studies. First, our analysis does not include considerations for staffing policies, as we do not have information on staffing. As mentioned previously, companies could consider alternate staffing policies to balance the working load across same- and next-day orders. One interesting research question could be to investigate the relationship between staffing level at a given hour and the behavior of individual agents (i.e., do workers' behavior differ when you work with 500 other colleagues in the same environment, as opposed to just 50?). Another fruitful avenue of research would be to examine whether there are internal or implicit deadlines for workers to meet at each point of delivery, and how these deadlines shape delivery performance outcomes and workers' behavior.

Another interesting avenue of follow-up study would be to consider other policies in place of the same-/next-day policy that JD.com currently has in place. For instance, JD.com could move the same-/next-day threshold to an earlier (e.g., 10 am) or later (e.g., 12 pm) time (by running staggered A/B tests), and observe how customers' behavior – as well as the behavior of individual workers – change as a result. Additionally, JD.com could also consider the "24-hour delivery" policy as described in the managerial implications above.

Finally, our study could benefit from more extensive inventory data. Additional details such as the number of SKUs available per day at each warehouse or inventory replenishment policies would enable researchers to study the role that inventory policies and stockouts play in same- and next-day deliveries. If stockouts are indeed one of the major factors that contribute to delay occurrences, the company could consider allowing customers to gain greater operational transparency into the inventory information, in case their orders take longer than their allotted time.

Overcoming Jargons in Consumer Contracts through Information Salience

MoonSoo Choi*
Harvard Business School, mochoi@hbs.edu

Monika Westphal*
Technion—Israel Institute of Technology, westphal@campus.technion.ac.il
* *These authors contributed equally*

Ryan Buell
Harvard Business School, rbuell@hbs.edu

Anat Rafaeli
Technion—Israel Institute of Technology, anatr@technion.ac.il

Galit B. Yom-Tov
Technion—Israel Institute of Technology, gality@technion.ac.il

Consumer contracts (e.g., terms and conditions) are often long and convoluted. Consumers rarely read and understand the contracts before signing, even when these contracts may impose various (e.g., financial, privacy) risks. This paper illustrates that leveraging information salience in contracts — translating legalese in consumer contracts into lay terms and highlighting the main points — improves consumer perceptions of contracts and firms. Across four experimental studies, consumers who were exposed to information salience exhibited higher understanding of the contract and higher trust in the company. In turn, they were more willing to sign contracts and reported higher perceived value of the service or product. However, some of these positive effects were attenuated when the risk associated with the contract and the service/product was perceived as high. We conclude with a general discussion on how managers can apply our research findings in practice.

1. Introduction

When signing up for a service, the burden of understanding the contract (also "terms" or "terms and conditions") falls on consumers. Regardless of how comprehensible consumer contracts are, consumers — rather than companies — are responsible for understanding and complying with the terms (Martin 2016). Nonetheless, consumers often fail to understand the contracts, as these are typically lengthy and include many unfamiliar terms (legalese)

(Marotta-Wurgler 2011, Ayres and Schwartz 2014, Bakos et al. 2014). In particular, studies have shown that anywhere between 74% (Obar and Oeldorf-Hirsch 2019) and 99.8% (Bakos et al. 2014) of consumers skip reading terms of services before agreeing. This "no-reading problem" has two main consequences. For one, consumers may take inadvertent (e.g., financial, privacy) risks by signing them. Additionally, some companies may take advantage of consumers' lack of understanding and include harmful, controversial provisions (Ayres and Schwartz 2014, Wilkinson-Ryan 2014). As of consequence, consumers may ultimately lose trust in the company and decide not to use its services in the future.

We propose a novel solution to consumers' failure to read and understand contracts, namely translating convoluted legal content into lay terms and outlining the main points of the contract. This form of *information salience* unveils the meaning of the contract created by extensive use of legal jargon and subsequently shortens the parts that should be read. We propose to provide summaries *alongside* the original version of the contract; as such, our design allows both consumers and companies to refer to the full contract as needed. We posit that our solution not only creates real opportunity for consumers to read contracts but also serves as a signal for positive company attributes. Indeed, the results suggest that making information salient in consumer contracts will increase consumers' understanding of the contract, as well as consumers' overall trust and positive perceptions of the company.

Across four experimental studies, we show that information salience increases consumer contract understanding and trust in the company, resulting in improved consumer willingness to sign contracts and perceived value of the service or product. Additionally, we identify contract risk as a boundary condition that attenuates the positive effect of information salience. Our findings have important theoretical implications for research, as well as practical implications for companies in improving consumer-firm relationships.

1.1. Problem Statement: Consumers Fail to Read and Understand Contracts

"Perhaps the most serious problem that deters reading [...] is the manner of presenting terms." – *Marotta-Wurgler (2011)*

Consumer contracts are often too long and contain unnecessary legal jargon for consumers to read and understand (Volokh 2008, Marotta-Wurgler 2011, Bakos et al. 2014, Furth-Matzkin and Sommers 2020). Even when consumers have sufficient time, consumers fail to read contracts, whether deliberately or inadvertently (Ayres and Schwartz 2014). A study with 48,154 visitors to various websites of software companies shows that less than 0.2% of the visitors accessed the user license agreement—and even those <0.2% who access the agreement just read a small portion (Bakos et al. 2014). Another study found that 74% of participants skipped reading the terms and conditions altogether, despite the contract requiring consumers to sell their firstborns in exchange for access to social network services (Obar and Oeldorf-Hirsch 2019). Amazingly, when an online gaming store chain *GameStation* revised its terms and conditions as part of an April Fool's Day joke to grant the firm a "...non transferable option to claim, for now and for ever more, their [users'] immortal soul", over 7,500 consumers still signed the terms and conditions, and a bare 12% of consumers noticed the clause (Bosker 2010). We can safely assume that—had consumers not overlooked these horrid clauses and realized these seemingly playful clauses were indeed binding—these consumers would not have signed these contracts.

1.2. Causes for the No-Reading Problem: Contract Characteristics and Psychological Factors

There are two main causes to the no-reading problem. First, there are *contract characteristics* that discourage consumers from reading contracts. Multiple scholars have noted the linguistic and legal complexity of typical consumer contracts (Sovern et al. 2014): "The degree of literacy required to comprehend the average disclosure form and key contract terms simply is not within reach of the majority of American adults." (White and Mansfield 2002) In addition to their legal complexity, contracts often tend to be very long. According to one estimate, an average American would have to spend approximately 250 hours to read all the digital contracts required for using routine online services (Berreby 2017). As these long contracts are loaded with legalese, an average consumer frequently fails to understand the implications of the text (Sovern et al. 2014). Thus, even though consumers can theoretically find and access all the information they need in the contracts, consumers' lack of reading and understanding these contracts keeps them from actually being informed (Bakos et al. 2014).

In addition to the contract characteristics, multiple *psychological factors* further contribute to the no-reading problem. First, people have limited information processing capability. People are often inundated by the overwhelming length and complexity of contracts (Eisenberg 1995). Second, people actively avoid seeking information when they sense a lack of control over negative outcomes (Sweeny et al. 2010). Consequently, consumers skip reading because they want to avoid feeling discomfort arising from uncertainty. At the same time, consumers likely become frustrated and annoyed, and decide not to read further (Wilkinson-Ryan 2014). And, even those who actually read the contracts pay close attention only to what seems most relevant and skip the rest (Milne and Culnan 2004). Third, people often underestimate what could go wrong in the future. Because people often do not appreciate future risks, they find many of the legal guardrail clauses irrelevant and consequently decide to skip reading altogether (Wilkinson-Ryan 2014). Fourth, consumers do not enjoy spending time assessing the benefits and risks of each clause in the contract (Hillman 2006), and tend to rely solely on their intuition and past experience (Eisenberg 1995). Consumers typically do not experience legal situations such as arbitration or assessing liquidated damages, and therefore may erroneously assume they will not experience them in the future either. Lastly, there is a social standard of not reading contracts (Stark and Choplin 2009); over time people learn that contracts are long, confusing, and unclear, which further reinforces this habit formation.

In sum, there are many contract characteristics and psychological factors that inhibit consumers from reading contracts in full length. We propose that information salience, which highlights key points in a simpler, more visually appealing form, has the potential to solve the no-reading problem. It helps mitigate psychological biases (e.g., habit formation, information overload) that preclude consumers from actually reading and understanding the contracts.

1.3. Risks for Consumers, Resulting from the No-Reading Problem

Even in an unlikely case that consumers actually read an entire contract, they may be deceived by fraudulent and unfair legal clauses that are camouflaged in convoluted legalese (Hillman 2006). As an example, the telecommunications marketing company Vertrue promoted free 30-day subscriptions to consumers but continued to furtively charge subscription fees thereafter (Furth-Matzkin and Sommers 2020). The details were included in the fine print but in an inconspicuous manner; as such, the company continued to collect monthly subscription fees from unknowing consumers who signed up for the "free" trial (Furth-Matzkin and Sommers (2020); also see *State v. Vertrue, Inc. (2013)*). This example clearly shows that unclear legal text can reduce the chances that consumers will recognize misleading or harmful terms. More problematically, even when consumers notice potential issues with the fine print, they rarely dispute the validity and fairness of the problematic terms (Bakos et al. 2014, Bosker 2010). Consumers who have an incentive to bring such disputes (e.g., poor quality of service, product defect) mostly do so only *after* a transaction is made (Becher and Zarsky 2008). In short, although various contract clauses—both misleading and non-misleading ones—remain known to the companies, only few consumers will understand and actively address them (Furth-Matzkin and Sommers 2020, Best 1981). Consumers' lack of incentives to read contracts, as well as the status quo of how consumer contracts are designed, allows many contracts to remain lengthy and unnecessarily complex (Stark and Choplin 2009, Korobkin 2003, Hillman 2006). Information salience, as proposed in this research, can help tackle the issue of companies potentially harmful terms from consumers.

1.4. Risks for Companies, Following the No-Reading Problem

Additionally, consumers' failure to read and understand contracts can adversely impact firms themselves. Although many consumers are shown to have steadfast belief that contracts are binding and enforced as written, legal research has shown that contracts that are (i) inconspicuous, (ii) contain material deception, or (iii) are notably inconsistent with the standard consumer contracts can be deemed as void (Furth-Matzkin and Sommers 2020, Silverman and Wilson 2016). For a notable instance, in the above mentioned legal case *State v. Vertrue, Inc. (2013)*, the court eventually ruled the company to pay $30 million dollars in restitution after deceptively billing the subscription fees to consumers for two years—despite consumers signing and agreeing to the contract. For another instance, a clause stipulating "...the company has no responsibility to notify [the consumer] of any changes..." was struck as void and ordered Safeway—the company and the defendant—to pay $42 million to its consumers (see *Rodman v. Safeway, Inc. (2014)*). Other similar legal cases include *Nguyen v. Barnes & Noble Inc. (2014)*, *Long v. Provide Commerce (2016)* and *In re Zappos.com, Inc. (2013)*, in which the courts concluded that mere appearance of the terms in the consumer contract does not guarantee enforceability. In fact, companies often deliberately refrain from enforcing many of the clauses written on the contract (Mann and Siebeneicher 2008) and have also guaranteed low-cost warranties to consumers beyond the fine print (Schwartz and Wilde 1983, Bebchuk and Posner 2005). Although it is unclear how much companies benefit from presenting inconspicuous or deceiving terms, the above findings imply that there are large costs associated with consumers raising legal claims. Thus, absent the compelling incentive for companies to make consumer contracts ambiguous and convoluted, companies may be able to reduce legal costs by making the contracts more accessible to consumers through information salience.

1.5. Previously Proposed Solutions to the No-Reading Problem in Contract Law

Under the current legal framework in contract law, the two parties (drafting and receiving the contract) have to affirm they had the opportunity to read a contract before signing, a concept called "the duty to read doctrine" (Ayres and Schwartz 2014). The law presumes that both parties understand and agree to their contract before signing it. This may be a workable solution for certain contracts. However, consumers sign contracts on an increasingly frequent basis – be it for purchasing a phone or computer, renting an apartment or office, going through a medical procedure, or on vacation. As a result, enforcing such duty-to-read doctrine for every consumer contract is not feasible given the rapidly growing number of contracts that people are required to sign (Ben-Shahar and Schneider 2011), given that even the shortest and relatively routine terms and conditions span a few thousand words. To respond to the duty to read doctrine, consumer protection law has attempted to induce firms to create a real opportunity for consumers to read contracts (Ayres and Schwartz 2014). However, this still is not sufficient to prevent companies from making contracts lengthy and full of legalese, which discourage consumers from reading altogether.

Contract law also encourages companies to use clickwrap (i.e., forcing the consumer to actively access the contract and click on an "I agree" icon) instead of browsewrap (i.e., merely including the contract somewhere on the website) to present the contract to consumers (Dickens 2007). Courts have ruled that contracts presented via browsewrap were unenforceable, since browsewrap method does not give sufficient notice of the contract. Nonetheless, even the clickwraps do not necessarily solve the no-reading problem; in a study with 47,399 households, consumers were only 0.36% more likely to view end user licence agreements presented as clickwraps than as browsewraps (Marotta-Wurgler 2011). Regulators have attempted to subject companies to mandated disclosure in response to the no-reading problem, be it by requiring contracts to be more comprehensible (Korobkin 2003) or by highlighting important keywords. (Ben-Shahar and Schneider 2011). Although such policies help consumers who *actually read* contracts to understand them in greater depth, the vast majority of consumers still skip reading contracts entirely, leaving the no-reading problem unresolved (Marotta-Wurgler 2011, Ben-Shahar and Schneider 2011). We anticipate that our information salience solution will help increase the contract readership, as well as consumers' understanding of contracts, through presenting the information in a more readable fashion.

1.6. Research Contribution

Our paper shows that information salience can be a useful tool to create a real opportunity for consumers to read contracts, as information salience not only makes contracts easier to read but also signals the company's willingness to be candid with consumers about their policies. We show that highlighting the main summary points of the contract and breaking down complex legalese increase consumer contract understanding.

Our findings have important theoretical implications for research in information salience, firm operations, and consumer management. First, we show that it is possible to highlight unappealing information to consumers yet still induce positive consumer response. A large body of research in information salience has shown that unattractive information (e.g., anti-smoking ads, mandatory calorie posting, negative customer review) keeps consumers informed but suppresses consumer demand (Bollinger et al. 2011, Blake et al. 2018, Chetty et al. 2009, Noar et al. 2017). However, we find that information salience can help increase positive consumer perceptions (i.e., improved value perceptions, higher willingness to pay) through increasing consumer understanding and trust. We explore information salience in the context of consumer contracts in greater detail. We additionally explore the effect of heightened contract risk as a boundary condition.

Moreover, our findings have important managerial and policy implications. Results show that highlighting the key points of contracts in simple, clear terms and displaying them alongside the original contracts increases both consumer trust and consumers' understanding of the contracts. Based on our research, our information salience solution could be used by regulators and legislators to solve the no-reading problem. Information salience would benefit firms as well, as it increases consumers' trust and positive perceptions toward firms and their products.

2. Theory Development

2.1. Information Salience as a Positive Signal

Although companies know the true quality of their products, it remains largely unobservable to consumers (Connelly et al. 2011). Companies use cues such as low prices, advertising, and warranties to signal high quality to consumers (Boulding and Kirmani 1993). Companies do so in the hopes of increasing positive consumer perceptions of their products and consumers' willingness to purchase.

Such firm-consumer interactions are extensively studied in signaling theory. Signaling theory aims to characterize dynamics between parties who have unequal access to information, as well as to reduce the information asymmetry between those parties (Spence 1973, 2002, Debo 2010, Schmidt and Buell 2017). The signaler (insider; in our case, firms) must choose how to communicate the signaling information to the signal receiver (outsider; in our case, consumers). The signal may induce a certain change in the receiver's behavior, depending on how the signal is interpreted (Connelly et al. 2011). Insiders have access to a mix of positive and negative information, and they must decide how to communicate this information to outsiders.

Because signals may be incomplete or misleading (Atkinson and Rosenthal 2014), it is imperative that consumers find the signals both useful and trustworthy (Boulding and Kirmani 1993). In other words, when consumers find signals to be less trustworthy, consumers may either rely less on the signal, or form negative perceptions toward the company altogether. As an example, when consumers have low trust and high doubt in company's certain practices such as privacy, assuring consumers that they are protected from privacy issues (e.g., "we protect your data for you") paradoxically heightens consumers' concerns about privacy (Brough et al. 2020).

Highlighting the key aspects of the contract may not only ease the difficulty of reading, but also provides clarity into what is actually written in the contracts. By resolving such sources of mistrust, information salience serves as a signal of trustworthiness to consumers, thereby reducing information asymmetry between the company and its consumers. Thus, we anticipate that both the company (signalers) and consumers (receivers) benefit from information salience; consumers will have an increased understanding of the contract and the associated risks, while the company will benefit from enhanced consumer trust, as well as from increased consumer willingness to sign contracts and improved value perceptions of the company's service or product.

2.2. Effect of Information Salience on Consumer Understanding of the Contract

Previous research has shown that consumers seldom have a substantial understanding of contracts. For instance, contracts usually include arbitration clauses in *italics* or ALLCAPS for additional emphasis. However, only 43% of people recognize when contracts include such clauses, and only 9% further understand that the clause would prevent them from proceeding in court (Sovern et al. 2014).

Despite firms' efforts to induce consumers to pay greater attention to contracts, many consumers still misunderstand or confuse important elements of contracts (Sovern et al. 2014, Furth-Matzkin 2017). For example, consumers do not adequately understand the relationship between monthly payments and interest rates (Shu 2010). To help consumers understand the products better, some companies have provided visual cues. For instance, companies have displayed graphs depicting statistics on loan financing and repayment to increase the borrowers' financial well-being (Bertrand et al. 2010). Similarly, companies have also shown simpler prototypes of mortgage loan disclosures to aid consumers in understanding mortgage loans better (Lacko and Pappalardo 2004).And, when eye-tracking technology was used to identify and highlight portions of the contracts that consumers skipped reading, both the time spent reading and the understanding of the terms and conditions increased significantly (Steinfeld 2016).

These examples above illustrate the potential effectiveness of highlighting important information in helping consumers understand contracts. Our information salience solution seeks to build on these examples. We anticipate that translating camouflaged content into lay terms and outlining the main points of the contract will help consumers to read and better understand the contracts.

2.3. Effect of Information Salience on Consumer Trust in the Company

Consumer trust—defined as the willingness to accept vulnerability in the actions of the company (Martin 2016)—has been recognized as an important concept in consumer research (Martin 2018, Milne and Culnan 2004). Motivating trust and engagement is a crucial operational challenge (Buell et al. 2020), and companies have engaged in trust-building marketing strategies to enhance consumer engagement (Aguirre et al. 2015). Trust has become increasingly important especially in online settings, where the amount of information asymmetries and hence consumers' uncertainty is high (Stewart 2003). Companies seek to increase consumer trust through cues that signal their benevolence and reliability (Kim and Kim 2011).

A large body of psychology literature suggests that disclosing information—especially that of sensitive or vulnerable nature—can instill a sense of trust (Derlega 1984, Laurenceau et al. 1998, Buell and Norton 2011). For instance, research has shown that companies can increase consumer trust and willingness to purchase by voluntarily revealing the hidden costs associated with producing goods (Mohan et al. 2020). For another instance, disclosing the prices of competitors' products also has been shown to increase consumer trust, even when the revealed information is potentially damaging in acquiring prospective consumers (Trifts and Häubl 2003).

Relatively little research has been done on consumer trust specifically in the context of consumer contracts. However, previous research appears to hint that consumers have diminished trust in the contracts they read and sign. Some companies have taken advantage by including fraudulent yet unenforceable terms throughout the long and convoluted contracts (Ayres and Schwartz 2014). Although consumers often do not consult contracts until *after* an incident has occurred, consumers ultimately notice that many contracts are problematic, which substantially decreases their trust in the contracts themselves (Martin 2018, Furth-Matzkin 2018). Such diminished consumer trust, in turn, diminishes the credibility of signals or messages from the company—even when those signals are positive and intended to bolster consumers' trust (Brough et al. 2020).

These findings suggest that—in order for information salience to be a positive signal—any information made salient in the consumer contract must tackle the issue of consumers' lack of trust in consumer contracts. Consumer mistrust primarily emanates from two sources in the context of consumer contracts: (a) contracts are long and difficult to read, and, consequently, (b) there may lurk deceptive or harmful terms camouflaged throughout the contracts, even though most consumers skip reading entirely.

As such, providing greater clarity into contracts can be perceived as the company's intention to build trustworthy relationships with consumers (Eisingerich and Bell 2008). When a company puts visible effort into gaining consumer trust, consumers reciprocate the company's efforts (Urban 2004, Eisingerich and Bell 2008) and are more likely to engage with it (Rust and Verhoef 2005). Building trust has been key in reducing consumers' privacy concerns and improving relationships between consumers and businesses (Milne et al. 2008). These examples suggest the positive role of making firms' policies more salient on building trust. Thus, we propose that information salience will increase consumers' trust in the company, in addition to increasing consumers' understanding of the contract.

2.4. Effect of Information Salience on Consumer Willingness to Sign Contracts and on Consumer Value Perceptions

Making information salient can nudge consumers to shift their perceptions and behavior (Bollinger et al. 2011, Donnelly et al. 2018, Noar et al. 2017, Mohan et al. 2020, Chetty et al. 2009, Blake et al. 2018). For example, mandatory calorie posting for foods in restaurants made consumers purchase foods with less calories, the average calories per transaction fell by 6% (Bollinger et al. 2011). Graphic warning labels reduced the consumption of sugary drinks, by 15%, heightening negative affect and prompting people to consider health consequences (Donnelly et al. 2018). Showing both the benefits and tradeoffs of credit cards ("tradeoff-transparency") increased average monthly spending by 19% and defection after nine months by 33% (Buell and Choi 2021).

Company signals, such as privacy protection seals (e.g., BBBOnline or TRUSTe), increase consumers' willingness to reveal information, and promote positive perceptions about the organization (Miyazaki and Krishnamurthy 2002). Interface appearance, intrusiveness of the questions asked, and the display of a privacy policy also can influence the extent to which consumers are willing to provide information (John et al. 2011). A surprising finding, however, reveals that even contextual cues totally unrelated to the objective danger of information disclosure may affect peoples' willingness to disclose information (John et al. 2011).

Research on operational transparency—which aims to reveal processes that are typically hidden from consumers' views—shows that disclosing behind-the-work processes (e.g., showing a list of airlines being searched on a flight search engine) helps consumers to recognize and appreciate the firm's effort to serve them (Buell and Norton 2011). These increased perceptions of effort induced consumers' feelings of reciprocity, which—in turn—increased consumers' satisfaction and perceived value of the service. Similarly, allowing consumers to observe the employees preparing meals in a restaurant increased their quality perceptions by 22% (Buell et al. 2017).

Although operational transparency differs from information salience in that it unveils full processes previously hidden from consumers (versus increasing the prominence of — and directing the attention to — certain parts of the process), they both aim to increase the visibility of operational processes. Similarly, we hypothesize that information salience in consumer contracts will also help consumers appreciate the effort and subsequently increase consumers' value perception of the firm and willingness to sign contracts.

Hence, we draw from above research to construct the following set of hypotheses:

Hypothesis 1 *Information salience will improve consumers' (a) willingness to sign contracts, and (b) perceived value of the service (or product).*

In examining the effect of information salience on consumer outcomes (e.g. purchase interest, engagement), research on operational transparency has illustrated a positive impact of disclosing hidden operating processes on consumers' perceived effort and appreciation (Buell et al. 2017, Buell and Choi 2021), and trust (Mohan et al. 2020, Buell et al. 2020). We propose a new mechanism that explains the effect of information salience on consumer outcomes: consumer contract understanding. Specifically, we hypothesize that—to the extent that information salience induces higher levels of contract understanding—contract understanding promotes an increased willingness to engage with the company.

Hypothesis 2 *Consumers' contract understanding will mediate the effect of information salience on consumers' (a) willingness to sign contracts, and (b) perceived value of the service (or product).*

In addition, creating trust has been identified as an approach for companies to increase consumers' information disclosure (Milne and Culnan 2004, Wu et al. 2012) and perceived value (Mohan et al. 2020, Buell et al. 2020). Hence, we hypothesize that consumer trust in the company will mediate the effect of information salience on consumer outcomes.

Hypothesis 3 *Consumers' trust in the company will mediate the effect of information salience on consumers' (a) willingness to sign contracts, and (b) perceived value of the service (or product).*

2.5. Moderating Effect of Consumer Contract Risk

Consumers avoid engaging with—or purchasing products from—companies with high perceived risk (Hoffman et al. 1999). High perceived risk may not only diminish consumers' trust (Olivero and Lunt 2004), but also impair their quality perceptions (Buell and Norton 2011, Groth and Gilliland 2006). Moreover, consumers are more willing to disclose sensitive information when they believe others have done so (Acquisti et al. 2013), as doing so signals low risk (of disclosure). In contrast, when perceived risk is high, consumers believe that they have less control over the information disclosed (Goodwin 1991). This decreases their overall trust in the company and their willingness to disclose information (Goodwin 1991, Olivero and Lunt 2004). Thus, scholars agree that privacy policies (and other types of contracts) should be designed to assure consumers that disclosing their personal information involves low risk (Wu et al. 2012, Milne and Culnan 2004). Based on these findings, we propose that:

Hypothesis 4 *If contract risk is high, information salience will decrease consumers' (a) willingness to sign contracts, (b) perceived value of the service/product, as well as consumers' (c) contract understanding and (d) trust in the company.*

Figure 3.1. **Research Model**

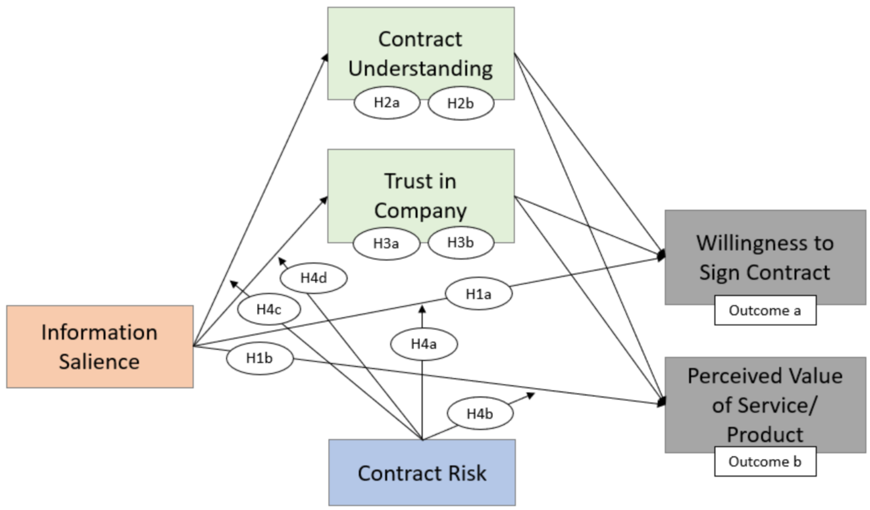

3. Presentation of Experiments

In four experimental studies, we test for the effect of information salience on consumers' willingness to sign contracts and perceived value of the service (or product) (see Figure 3.1. for the research model). We test our model in two different contexts: reading terms and conditions prior to using a social media website (Study 1–3), and reading a lease agreement prior to leasing an apartment (Study 4). We chose the two (very different) contexts because the apartment leasing context typically poses greater risk than using social media (e.g., risk of damaging the premises) and usually requires more personal attention to the contract. Both contexts represent typical scenarios of consumer experiences.

For an overview of all hypotheses, see Figure 3.1. We tested Hypothesis 1 (main effects) in Study 1. Hypothesis 2 (mediation of contract understanding) and Hypothesis 3 (mediation of trust in the company) were tested in Study 2–4. In Study 3, we tested the full moderated mediation model and therefore also included Hypothesis 4 (moderation of contract risk). In Study 4, we aimed to replicate the full moderated mediation model in another context than social media, i.e. leasing. In the following section, we will outline method, results and discussion for each of the studies separately. In sum, we find very strong evidence for Hypotheses

1–3 (main and mediating effects), but weak or no evidence for Hypothesis 4 (moderating effects).

In all experiments we pre-set our sample sizes to 100 per experimental condition and did not analyze the data until the target was met. We note that we set the number of participants to be roughly 10% higher than our target in all studies to ensure there were enough responses in each condition to draw meaningful conclusions. We report all manipulations, measures, and data exclusions.

4. Study 1

4.1. Procedure

In the beginning of the study, participants were asked to assume the role of an online user and imagine the following scenario: "Imagine that you are signing up for an account at the social media & cloud storage platform Coquiptur, which allows you to encrypt and secure your personal data using random images. In compliance with standard legal practice, Coquiptur requires all users to read and agree to its Terms and Conditions." Then, we showed all participants the terms of the fictitious company, Coquiptur, for this study (5,232 words—a reasonable length for such terms (Obar and Oeldorf-Hirsch 2018). Participants were randomly assigned to one of two information salience conditions; that is, the participants either saw the regular terms (condition 1: no summaries, see Figure 3.9), or they saw the terms with side-by-side summaries (condition 2: summaries, see Figure 3.10). Following the terms, we asked the participants a number of questions to test their understanding of the terms. Additionally, we asked the participants how much they agree with the terms, as well as their perceived value of the company's service. Lastly, we asked participants for their demographic information and their general perceptions of terms and conditions. We thanked them for their participation and paid them.

4.2. Participants

For the study, we recruited 212 participants online on Mechanical Turk. (*Mean* age=40 years, *SD* age=11.55, 53% female). Participants received $2 for participation in an 8-min study on consumers' perceptions of consent. One participant was excluded from the analysis as s/he did not pass the attention check. Hence, final sample size was 211. We use the terms 'contract' and 'terms and conditions' interchangeably in Study 1, 2 and 3.

4.3. Measures

The main measure of interest, *information salience*, was manipulated by displaying short side-by-side summaries in simple language alongside the full-length terms. No summaries were shown to the participants in the control condition. Hence, they only saw the full-length terms as consumers typically do.

We use the following two variables as performance measures:

Willingness to sign contracts. To measure willingness to sign contracts, we asked: "How much do you agree with the Terms and Conditions?" on a scale from '1 – not at all' to '7 – totally'.

Perceived value of the company's service was assessed along four dimensions: quality, emotional, price, and social (adopted from Sweeney and Soutar 2001). For each dimension we asked one question, measured on a scale from '1 – I totally disagree' to '7 – I totally agree': "Coquiptur's service is of high quality"; "Coquiptur's service is a service I would want (to use)"; "I would be willing to pay a lot for Coquiptur's service"; "Other people would approve of Coquiptur's service". Cronbach's Alpha was 0.848.

Other variables of interest included:

Time spent reading the contract, measured by the time (in minutes) participants spent on the page displaying the terms and conditions.

Contract understanding[1] We measured contract understanding by the following six items: "Coquiptur is the owner of the content I post" (false); "Coquiptur requires me to provide the login information to other third-party social platforms I link my account with" (true); "Coquiptur is responsible for any slowdown of the website" (false); "Coquiptur uses my data to improve the performance of the platform" (true); "Coquitur regularly backs up and stores user data" (true) and "Legal matters with Coquiptur are governed under the law of California" (true). We counted the number of questions answered correctly (ranging from 0 to 6) and used this score as a measure of subjects' understanding of the terms and conditions.

Reading habits. Participants reported their reading habits regarding terms and conditions (e.g., "It is normal to sign up for websites/apps without reading the Terms and Conditions")

Demographics. Finally, participants reported their demographics (i.e., gender, age, English proficiency, education, ethnicity).

4.4. Results

4.4.1. Descriptive analysis.

Correlations. For an overview of the correlations, see Table 3.5.

Controls. The amount of time spent reading the contract was significantly higher in the information salience condition ($Mean = 211.57, SD = 167.39$ vs. $Mean = 138.42, SD = 177.55$, $F(1, 209) = 9.480, p = 0.002$). In addition, subjects in the information salience condition exhibited significantly higher contract understanding ($Mean = 4.88, SD = 0.99$ vs. $Mean = 4.24, SD = 1.89$), $F(1, 209) = 23.829, p < 0.001$).

[1] At this stage, we added it as an exploratory but focal variable. In Study 2, we include it as a mediator.

Figure 3.2. Study 1: The Effect of Information Salience on Consumer Outcomes

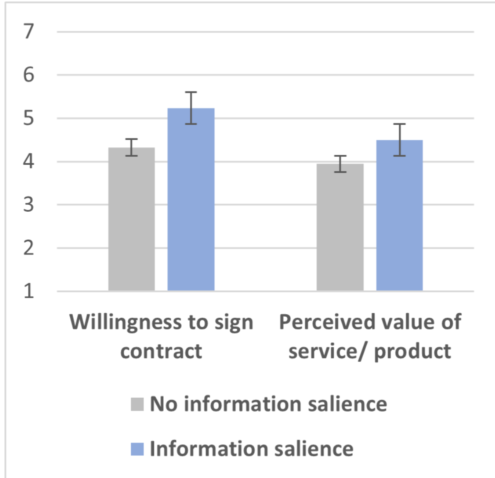

4.4.2. Hypothesis testing. We ran two linear regressions with control variables to test the effect of information salience on the two consumer outcomes: willingness to sign con-tracts and perceived value. All results are shown in Table 3.1. In both analyses, we controlled for participants' time spent reading the contract, contract understanding, contract reading habits, as well as their demographics (i.e., gender, age, English proficiency, education and ethnicity).

Hypothesis 1a. Information salience significantly increased subjects' willingness to sign contracts $(Mean = 5.23, SD = 1.41$ vs. $Mean = 4.32, SD = 1.57, p < 0.001)$, $F(9,201) = 4.558, R^2 = 0.170, p < 0.001$, see Figure 3.2. Hence, Hypothesis 1a is supported.

Hypothesis 1b. Information salience also significantly increased subjects' perceived value of the company's service $(Mean = 4.50, SD = 1.17$ vs. $Mean = 3.94, SD = 1.24, p < 0.001)$, $F(9,201) = 4.641, R^2 = 0.172, p < 0.001$ (see Figure 3.2.). Hence, Hypothesis 1b is also supported.

4.5. Discussion

The results of Study 1 suggest that consumers appreciate information salience. In particular, consumers who saw summaries of the contract spent more time reading the terms and exhib-ited better understanding of the contract. Also, they were more likely to sign contracts and

Table 3.1 Regression Results Study 1—Willingness to Sign and Perceived Value of Service

	Willing. to Sign		**Perc. Value**	
	coeff	*se*	*coeff*	*se*
Constant	5.33^{***}	1.42	5.77^{***}	1.10
Information salience	0.94^{***}	0.21	0.59^{***}	0.16
Reading habits	$-0.23^{\dagger}$	0.12	-0.34^{***}	0.09
Gender	-0.46^{*}	0.21	$-0.30^{\dagger}$	0.16
Age	-0.01	0.01	-0.02^{*}	0.01
Language	0.40	1.07	0.64	0.83
Education	-0.04	0.11	0.00	0.08
Ethnicity	-0.08	0.08	-0.04	0.06
N	211		211	
R^2	0.137		0.162	
Adj. R^2	0.107		0.133	
Residual Std. Error	1.480		0.133	
F Statistic (df)	4.591^{***} (7,203)		5.601^{***} (7,203)	

Note. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; † $p < 0.10$

reported higher perceived value of the company's service. Our initial results suggest positive effects of providing information salience in consumer contracts. In the next study, we test whether consumer contract understanding and trust in the company mediate the positive effect of information salience on consumers' willingness to sign contracts and perceived value of the service.

5. Study 2

5.1. Participants and Procedure

For Study 2, we ran another online study on Mechanical Turk that builds on the previous study. The study procedure was identical to the procedure used for Study 1, except that we also asked participants for their *perceived* understanding of the terms and their trust in the company. The study included 222 participants ($Mean$ age=38 years, SD age=11.58, 46% female) that received $2 for participating in an 8-min study.

5.2. Additional Measures

In addition to the variables described in Section 4.3, we measured the following two variables that were included as mediators (besides contract understanding):

Perceived contract understanding. We adapted and used the three items from Obar and Oeldorf-Hirsch (2019), on a scale from '1 – I totally disagree' to '7 – I totally agree': "The language in the Terms and Conditions is clear"; "The Terms and Conditions are difficult to understand"; "The Terms and Conditions provide helpful information". Cronbach's Alpha was 0.924.

Trust in the company was measured by means of three items on a scale from '1 – I totally disagree' to '7 – I totally agree', as adapted from Martin and Murphy (2017): "Coquiptur is a trustworthy company"; "I have confidence in Coquiptur's behaviors"; "I think Coquiptur is a reliable company". Cronbach's Alpha was 0.950.

5.3. Results

5.3.1. Descriptive analysis.

Correlations. For an overview of the correlations, see Table 3.6.

Controls. Although subjects in the information salience condition spent about 20 seconds more on average reading the terms, information salience did not significantly increase time spent reading the contract ($Mean = 163.98, SD = 124.89$ vs. $Mean = 141.02, SD = 188.53$), $F(1, 220) = 1.171, p = 0.280$.

5.3.2. Hypothesis testing. To test for the mediating role of (1) contract understanding, (2) perceived contract understanding, and (3) trust in the company in the effect of information salience on consumer outcomes, we ran a mediation analysis (mediation model no. 4, Hayes 2017). We ran two models: one with willingness to sign contracts as the outcome—to test Hypotheses 2a and 3a, and one with perceived value as the outcome—to test Hypotheses 2b and 3b. All results appear in Table 3.2.

Both models include the three mediators and all control variables, i.e., time spent reading the contract, terms and conditions' reading habits, and demographics (i.e., gender, age, English proficiency, education and ethnicity). We report below step by step the results of the two mediation models, separately.

The Effect of information salience on willingness to sign contracts.

The full mediation model was significant, $F(11, 210) = 42.143, R^2 = 0.672, p < 0.001$.

Hypothesis 2a. Information salience significantly increased contract understanding ($Mean = 4.68$, $SD = 1.09$ vs. $Mean = 4.44$, $SD = .89$, $p = 0.054$) and perceived contract understanding ($Mean = 5.20$, $SD = 1.54$ vs. $Mean = 3.91$, $SD = 1.61$, $p < 0.001$), see Figure 3.3. However, neither contract understanding nor perceived contract understanding mediated the effect of information salience on willingness to sign contracts, regardless of trust in the company. Hence, Hypothesis 2a is not supported.

Hypothesis 3a. Information salience significantly increased trust in the company ($Mean$ $= 4.92$, $SD = 1.53$ vs. $Mean = 4.11$, $SD = 1.35$, $p < 0.001$), see Figure 3.3. Addition-ally, trust in the company mediated the effect of information salience on willingness to sign contracts, regardless of contract understanding and perceived contract understanding ($p <$ 0.001). Hence, Hypothesis 3a is supported.

The Effect of information salience on perceived value.

Figure 3.3 Study 2: The Effect of Information Salience on Consumer Outcomes

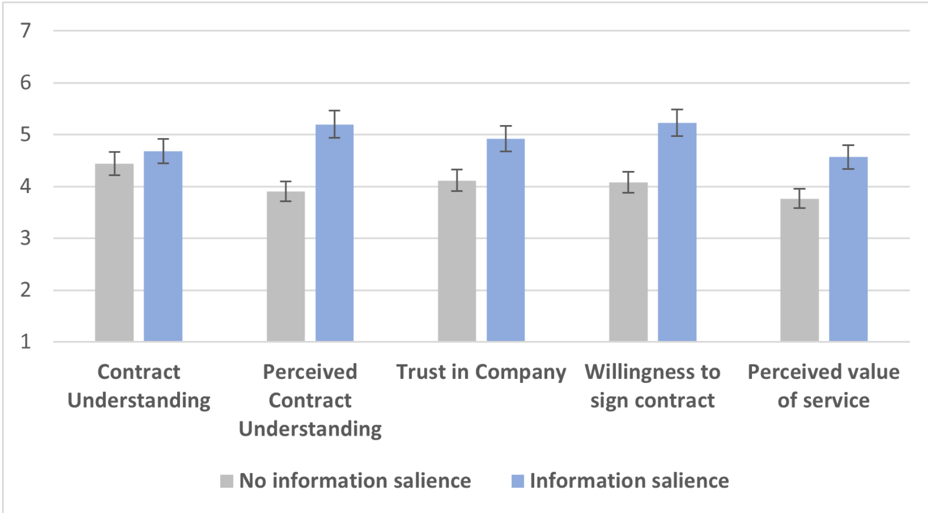

The full mediation model was significant, $F(11,210) = 36.108, R^2 = 0.626, p < 0.001$.

Hypothesis 2b. Information salience significantly increased contract understanding ($Mean = 4.68, SD = 1.09$ vs. $Mean = 4.44, SD = 0.89, p = 0.054$) and perceived contract understanding ($Mean = 5.20, SD = 1.54$ vs. $Mean = 3.91, SD = 1.61, p < 0.001$). However, neither contract understanding nor perceived contract understanding mediated the effect of information salience on perceived value, regardless of trust in the company. Hence, Hypothesis 2b is not supported.

Hypothesis 3b. Information salience significantly increased trust in the company ($Mean = 4.92$, $SD = 1.53$ vs. $Mean = 4.11$, $SD = 1.35$, $p < 0.001$), see Figure 3.3. Furthermore, trust in the company mediated the effect of information salience on perceived value, regardless of contract understanding and perceived contract understanding ($p < 0.001$). Hence, Hypothesis 3b is supported.

5.4. Discussion

We find that Study 2 replicated the results of Study 1, and in doing so reaffirm the conclusions of Study 1. We again see that information salience increases time spent reading the contract and that it improves consumer outcomes (i.e., willingness to sign and perceived value of

Table 3.2 Mediation Results Study 2—Willingness to Sign and Perceived Value

	MEDIATORS						OUTCOME			
	Understand.		**P. Understand.**		**Trust**		**Willing. to Sign**		**Perc. Value**	
	coeff	*se*	*coeff*	*se*	*coeff*	*se*	*coeff*	*se*	*coeff*	*se*
Constant	3.01***	0.69	5.06***	1.09	5.28***	1.01	-0.03	0.67	0.59	0.63
Information salience	0.28*	0.14	1.27***	0.21	0.75***	0.20	0.46***	0.13	0.20†	0.12
Understanding							0.07	0.06	-0.07	0.06
Perc. understanding							0.08	0.06	0.08	0.05
Trust							0.71***	0.06	0.62***	0.06
Reading habits	0.23**	0.08	-0.21†	0.12	-0.16	0.11	0.00	0.07	-0.01	0.07
Gender	-0.14	0.13	-0.01	0.21	0.27	0.19	0.03	0.12	0.12	0.11
Age	0.00	0.01	-0.01	0.01	-0.01	0.01	-0.01	0.01	0.00	0.00
Language	-0.27	0.30	0.19	0.47	0.15	0.43	-0.03	0.26	0.08	0.25
Education	0.01	0.07	-0.25†	0.11	-0.23*	0.10	0.08	0.06	0.06	0.06
Ethnicity	0.01	0.04	-0.02	0.07	-0.02	0.06	0.04	0.04	0.05	0.03
N	222		222		222		222		222	
R^2	0.068		0.191		0.117		0.688		0.643	
Adj. R^2	0.033		0.161		0.083		0.672		0.624	
Residual Std. Error	0.978		2.405		2.060		0.761		0.662	
F Statistic (df)	2.232* (7,214)		7.230*** (7,214)		4.061*** (7,214)		46.577*** (10,211)		38.0212*** (10,211)	

Note. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; † $p < 0.10$

the service). In addition, this study shows that consumers'—both perceived and actual—contract understanding and trust in the company improve and explain the mechanism for the improvement in consumer outcomes. Next, we test whether the presence of high contract risk attenuates the positive effect of information salience on consumer outcomes.

6. Study 3

6.1. Participants and Procedure

We ran another online study on Mechanical Turk ($N = 456$), similar to the previous studies. Participants (*Mean* age=40 years, *SD* age=12.19, 53% female) received $2 for participating in an 8-min study. Two participants failed the attention check and were thus excluded from the analysis. In addition, 40 participants (8.8%) failed the manipulation check for contract risk, and were thus excluded from the analysis. Our final sample had 414 participants (91% of original sample).

In addition to replicating the main effect of information salience and the mediating effects of contract understanding and trust in the company, we explore the moderating role of risk in contracts. We examine both in the (direct) effect of information salience on consumers' willingness to sign contracts and perceived value of the service, as well as in the (indirect) effect on contract understanding and trust in the company. Hence, we used a 2 (information salience: summary vs. no summary) × 2 (contract risk: high vs. low) between subjects design, where participants were randomly assigned to either a low or a high contract risk condition, regardless of whether they saw summaries of the terms or not. Participants were again informed that "in compliance with legal practice, Coquiptur requires all users to read and agree to its Terms and Conditions". Additionally, we also informed them: "note that Coquiptur will use your data—which is **anonymized** (for low risk condition) vs. **unanonymized** (for high risk condition) and **cannot be (low risk)** vs. **can be (high risk)** traced back to you—in the future" (see Figures 3.13 and 3.14). We also embedded the sentence in the terms themselves, see Figures 3.11 and 3.12. A pretest confirmed that the (privacy) risk was indeed perceived as low/high in these two conditions. We added a manipulation check for contract risk to the end of the study.

6.2. Additional Measures

Manipulation check Contract Risk. We ran a manipulation check to see if people paid attention to the contract risk condition they were in, using the following items: "The Terms and Conditions mentioned that Coquiptur would use my data, ... (a) but it could not be traced back to me (anonymized data)"; (b) and it could be traced back to me (unanonymized data)"; (c) I don't remember". We excluded participants who did not correctly remember their assigned contract risk conditions.

6.3. Results

6.3.1. Descriptive analysis.

Correlations. For an overview of the correlations, see Table 3.7.

Controls. We notice that people in general spent much more time reading the terms, compared to the previous two studies. Though people spend about 20 seconds more on average reading the terms, information salience did not significantly increase time spent reading the contract ($Mean = 236.00, SD = 287.99$ vs. $Mean = 213.66, SD = 255.09$), $F(1,412) = 0.694, p = 0.405$.

6.3.2. Hypothesis testing. To test for the mediating role of (1) contract understanding, (2) perceived contract understanding, and (3) trust in the company, as well as the moderating role of contract risk in the effect of information salience on consumer outcomes, we ran a moderated mediation analysis (model no. 8, Hayes 2017). We ran two models: one with willingness to sign contracts as the outcome—to test Hypotheses 2a–4a, and the other with perceived value as the outcome—to test Hypotheses 2b–4b. All results appear in Table 3.3. Both models included estimates to test Hypotheses 4c and 4d (moderating effect of contract risk on mediators). We report below the results of the two moderated mediation models, separately, step by step. All analyses reported below include the three mediators, the moderator,

Figure 3.4 Study 3: The Effect of Information Salience on Consumer Outcomes

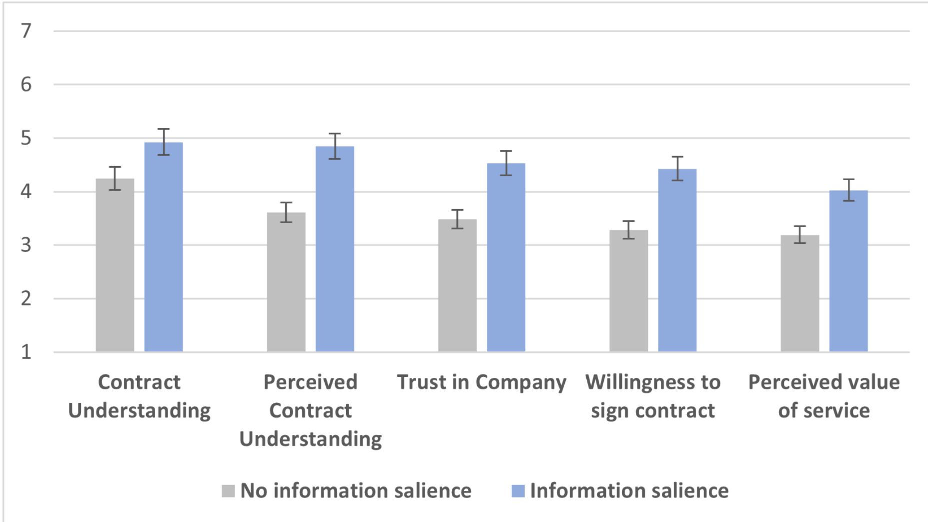

and all control variables, which include time spent reading the contract, general perceptions of terms and conditions, and demographics (i.e., gender, age, English proficiency, education and ethnicity).

The Effect of Information Salience and Contract Risk on Willingness to Sign Contracts.

The full moderated mediation model was significant, $F(13, 400) = 82.973, R^2 = 0.730, p < 0.001$.

Hypothesis 2a. Information salience significantly increased contract understanding ($Mean = 4.93, SD = 1.03$ vs. $Mean = 4.25, SD = 1.01, p < 0.001$), and perceived contract understanding ($Mean = 4.85$, $SD = 1.57$ vs. $Mean = 3.61$, $SD = 1.64$, $p = 0.027$), see Figure 3.4. And, both contract understanding and perceived contract understanding (marginally) mediated the effect of information salience on willingness to sign contracts ($p = 0.008$, $p = 0.058$, respectively), regardless of trust in the landlord. Hence, Hypothesis 2a is supported.

Hypothesis 3a. Information salience increased trust in the landlord ($Mean = 4.53$, $SD = 1.53$ vs. $Mean = 3.48$, $SD = 1.49$, $p < 0.001$). Furthermore, trust mediated the effect of information salience on willingness to sign contracts. Hence, again Hypothesis 3a is supported.

Table 3.3 Study 3: Moderated Mediation Results—Willingness to Sign and Perceived Value

	MEDIATORS						OUTCOME			
	Understand.		P. Understand.		Trust		Willing. to Sign		Perc. Value	
	coeff	*se*	*coeff*	*se*	*coeff*	*se*	*coeff*	*se*	*coeff*	*se*
Constant	2.69***	0.60	5.33***	0.91	6.32***	0.85	0.69	0.55	1.72***	0.49
Information salience	0.84**	0.32	2.06***	0.49	1.63***	0.46	-0.11	0.28	-0.28	0.25
Understanding							0.12**	0.04	0.06	0.04
Perc. understanding							0.07†	0.03	0.07*	0.03
Trust							0.77***	0.04	0.60***	0.03
Risk	0.29*	0.14	0.17	0.22	-0.46*	0.21	-0.37**	0.13	-0.30**	0.11
Interaction (info x risk)	-0.07	0.20	-0.64*	0.31	-0.47	0.29	0.18	0.17	0.22	0.15
Reading habits	0.11†	0.06	-0.35***	0.09	-0.27**	0.09	0.01	0.05	-0.06	0.05
Gender	0.07	0.09	0.15	0.14	0.02	0.13	0.03	0.08	-0.05	0.07
Age	0.00	0.00	-0.01	0.01	-0.01	0.01	-0.01**	0.00	-0.01*	0.00
Language	0.18	0.26	0.66†	0.40	0.23	0.37	-0.17	0.22	-0.23	0.20
Education	0.02	0.05	-0.13	0.08	-0.12	0.08	0.05	0.05	0.07†	0.04
Ethnicity	-0.04	0.04	0.09	0.06	0.00	0.05	-0.01	0.03	0.03	0.03
N	414		414		414		414		414	
R^2	0.123		0.202		0.193		0.730		0.675	
Adj. R^2	0.108		0.182		0.173		0.721		0.664	
Residual Std. Error	1.024		2.407		2.111		0.749		0.597	
F Statistic (df)	6.690*** (9,404)		11.331*** (9,404)		10.703*** (9,404)		90.109*** (12,401)		69.294*** (12,401)	

Note. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; † $p < 0.10$

Hypothesis 4a. Contract risk did not moderate the effect of information salience on willingness to sign contracts ($p = 0.315$). Hence, Hypothesis 4a is not supported.

The Effect of information salience and Contract Risk on Perceived Value.

The full moderated mediation model was significant, $F(13, 400) = 64.047, R^2 = 0.676, p < 0.001$.

Figure 3.5 Study 3: The Effect of Contract Risk on Consumer Outcomes

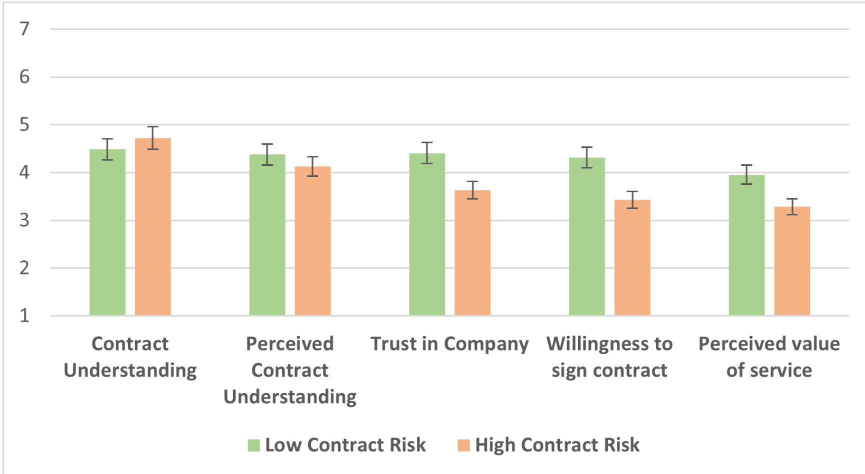

Hypothesis 2b. As shown above, information salience significantly increased contract understanding ($p < 0.001$) and perceived contract understanding ($p = 0.027$). On the other hand, Contract understanding did not mediate the effect of information salience on perceived value ($p = 0.183$). However, perceived contract understanding mediated the effect of information salience on perceived value ($p = 0.053$), regardless of information salience and trust in the company. Hence, Hypothesis 2b is partly supported.

Hypothesis 3b. We already know that information salience increased trust in the landlord ($p < 0.001$). Furthermore, trust mediated the effect of information salience on perceived value ($p < 0.001$), regardless of contract understanding. Hence, Hypothesis 3b is again supported.

Hypothesis 4b. Contract risk did not moderate the effect of information salience on perceived value ($p = 0.151$). Hence, Hypothesis 4b is not supported.

Hypothesis 4c. Contract risk did not moderate the effect of information salience on contract understanding ($p = 0.764$), but it moderated the effect of information salience on perceived contract understanding ($p = 0.040$), see Figure 3.6. Hence, Hypothesis 4c is partially supported.

Hypothesis 4d. Contract risk did not moderate the effect of information salience on trust in the company ($p = 0.101$). Hence, Hypothesis 4d is not supported.

Figure 3.6 The Moderating Effect of Contract Risk in the Effect of Information Salience on Consumer Perceived Contract Understanding (Study 3)

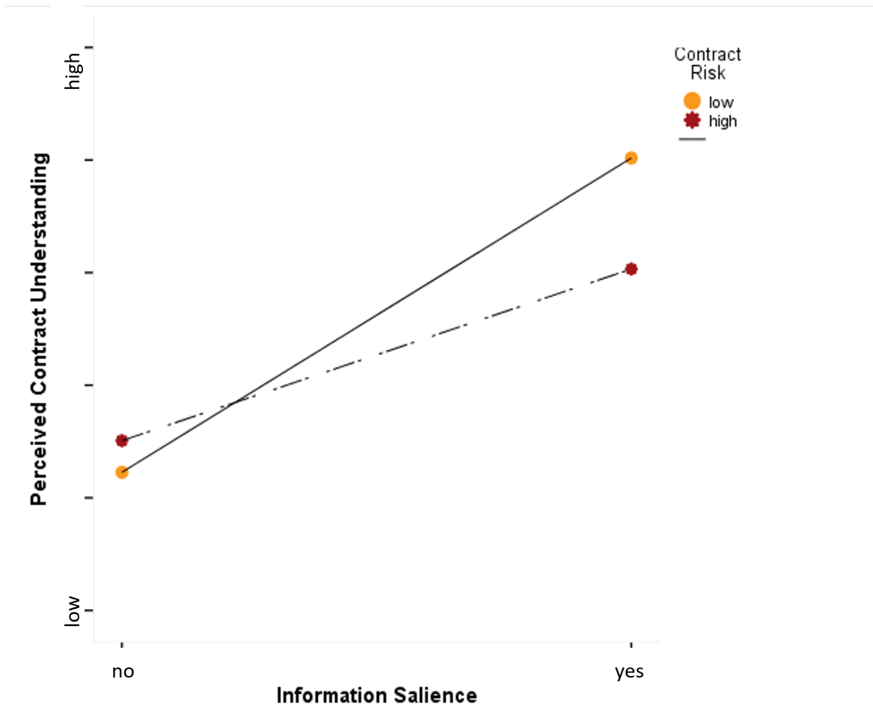

6.4. Discussion

First, we were able to replicate most of the results of Study 2 in Study 3. We again found evidence supporting the positive effect of information salience on consumer outcomes. In particular, consumers exposed to information salience exhibited higher contract understanding and trust in the company. Increased willingness to sign contracts and improved perceived value of the service were explained by both consumers' increased contract understanding and trust in the company.

Additionally, we also observed how consumer outcomes change in the presence of contract risk. In particular, higher contract risk was associated with an increase in actual contract understanding but with a decrease in consumers' perceptions of the contract and the company (i.e., perceived contract understanding, trust in the company, willingness to sign contracts, and perceived value). These results suggest that the positive effects of information salience may be attenuated by the presence of high contract risk, although the moderation effect was insignificant. These results suggest that, although consumers generally appreciate information salience, highlighting contract risks may cause consumers to be more disproportionately alert to the risks.

7. Study 4

7.1. Participants and Procedure

We ran another online study on Mechanical Turk (N=455). Participants ($Mean$ age=40 years, SD age=11.46, 51.3% female) received $2 for participation in an 8-min study on apartment leasing. One participant failed the attention check and was thus excluded from the analysis. Hence, final sample size was 454. The context of this study was reviewing a leasing contract; this context represents a scenario involving a higher and more tangible risk. We want to examine whether the findings of Study 3 replicate in a different, but still very prevalent context. Hence, we again test for the moderating role of contract risk (here: financial) in the (direct) effect of information salience on consumers' willingness to sign contracts and perceived value (of the property), as well as in the (indirect) effect on contract understanding and trust. Hence, we again used a 2 (information salience: summary vs. no summary) $\times$ 2 (contract risk: high vs. low) between-subjects design.

Study procedure was very similar to the previous studies, but in an apartment leasing context. Participants were provided with a different scenario: "Imagine you just moved to a new city and were happy to find a property that you will be staying in for 1 year, with the monthly rent of $1,000. In order for a prospective resident to lease the property, the resident needs to read and sign the lease agreement. The blanks in the lease agreement are already pre-filled." Participants were either assigned to the high or low contract risk condition (as part of the lease agreement scenario): "A breach of this lease agreement will result in a $500 (low risk) / \$5,000 (high risk) penalty fee" (see Figure 3.17 and Figure 3.18). A pretest confirmed that a breach of \$500 was perceived as low, and \$5,000 as high contract risk.

Next, we presented all participants with the lease agreement (4,796 words, shorter than the terms and conditions in the previous studies, but a reasonable length for a lease agreement). Participants were randomly assigned to one of two conditions. That is, they either saw the full, original version of the lease agreement (condition 1: no summary, see Figure 3.15), or they saw it with a side-by-side summary (condition 2: summary, see Figure 3.16).

Following the lease, we tested for participants' understanding of the lease agreement, and asked them for their perceived understanding of the terms and their trust in the landlord. Also, we asked them how likely they would sign the lease agreement, and how they perceive the value of the property. Lastly, participants reported their lease agreement reading habits and demographics. We thanked them for their participation and they received payment.

7.2. Measures

Contract risk was manipulated by the introduction of a penalty fee for any breach of the contract, i.e., \$500 for low risk, and \$5,000 for high risk. In particular, the sentence "A breach of this lease agreement will result in a \$500 / \$5,000 penalty fee" was added to both the introduction and the beginning of the lease agreement.

Contract understanding was measured by means of six items: "When natural disasters happen, I can end the lease agreement"; "The landlord is responsible for my personal property (e.g., car)" (false; "Foster children and live-in aides need the landlord's approval to live in the property" (true); "If there is a legally invalid clause in the lease, all remaining parts of the lease agreement would be invalid" (false); "My failure to pay the rent would immediately terminate the lease agreement" (true); "Criminal activity away from the property would result in lease termination" (true) and "I am required to let the landlord in the property for any repairs" (true).

Perceived contract understanding was measured as based on (Obar and Oeldorf-Hirsch 2019): "The language in the lease agreement is clear"; "The lease agreement is difficult to understand"; "The lease agreement provides helpful information". Cronbach's Alpha was 0.848.

Trust in landlord was assessed following Martin and Murphy (2017): "I trust the landlord"; "The landlord is very trustworthy"; "I have confidence in the landlord's behaviors"; "The landlord is reliable". Cronbach's Alpha was 0.947.

Willingness to sign contracts. We asked: "How likely would you be to sign the lease agreement?", on a scale from '1 - Not at all likely' to '7 - totally likely.

Perceived value were assessed as in the previous studies, but the items were adapted to the new context: "The property is of high quality"; "The property is a property I would want"; "I would be willing to pay a lot for the property"; "Other people would approve of the property". Cronbach's Alpha was 0.871.

Manipulation check. We ran a manipulation check to see if people indeed perceived the contract risk as high vs. low. In particular, participants responded to the question "Please indicate if you consider the penalty for breaching the terms in the lease agreement as low or high for each of the following two amounts: $500 and $5,000" (on a scale of '1 - very low' to '7 - very high'). The two amounts were presented to participants in random order.

Control variables. We measured the same variables as in the previous studies. Participants additionally reported relationship status and household income, as we considered these variables relevant to the study context.

7.3. Results

7.3.1. Descriptives.

Correlations. For an overview of the correlations, see Table 3.8.

Controls. Information salience did not significantly increase time spent reading the contract ($Mean = 296.69, SD = 381.09$ vs. $Mean = 286.50, SD = 288.87, p = 0.48$).

7.3.2. Hypothesis Testing To test for the mediating role of (1) contract understanding,(2) perceived contract understanding, and (3) trust in the landlord, and the moderating role of contract risk on the effect of information salience on consumer outcomes, we ran a moderated mediation analysis (model no. 8, Hayes 2017). We ran the model twice. Once with willingness to sign contracts as the outcome to test Hypotheses 1a–4a, and once with perceived value as the outcome to test Hypotheses 1b–4b. All results appear in Table 3.4. Both models also included estimates to test Hypotheses 4c and 4c (moderating effect of contract risk on mediators). We report below step by step the results of the two moderated mediation models, separately. All analyses reported below include: the three mediators, the moderator, and all control variables: time spent reading the contract, terms and conditions' reading habits, and demographics (i.e., gender, age, English proficiency, education, ethnicity, relationship status and household income).

The Effect of Information Salience and Contract Risk on willingness to Sign Contracts.

The full moderated mediation model was significant, $F(15, 438) = 21.648$, $R^2 = 0.426$, $p < 0.001$.

Hypothesis 2a. Information salience significantly increased contract understanding ($Mean = 4.66$, $SD = 1.01$ vs. $Mean = 4.25$, $SD = 1.05$, $p < 0.001$), and perceived contract understanding ($Mean = 5.16$, $SD = 1.33$ vs. $Mean = 4.72$, $SD = 1.51$, $p = 0.027$), see Figure 3.7. In addition, both contract understanding and perceived contract understanding mediated the effect of information salience on willingness to sign contracts ($p = 0.072$, $p = 0.003$), regardless of trust in the landlord. Hence, Hypothesis 2a is supported.

Hypothesis 3a. Information salience did not increase trust in the landlord ($Mean = 4.89, SD = 1.28$ vs. $Mean = 4.72, SD = 1.35, p = 0.138$), see Figure 3.7. Hence, though trust in

Figure 3.7 Study 4: The Effect of Information Salience on Consumer Outcomes

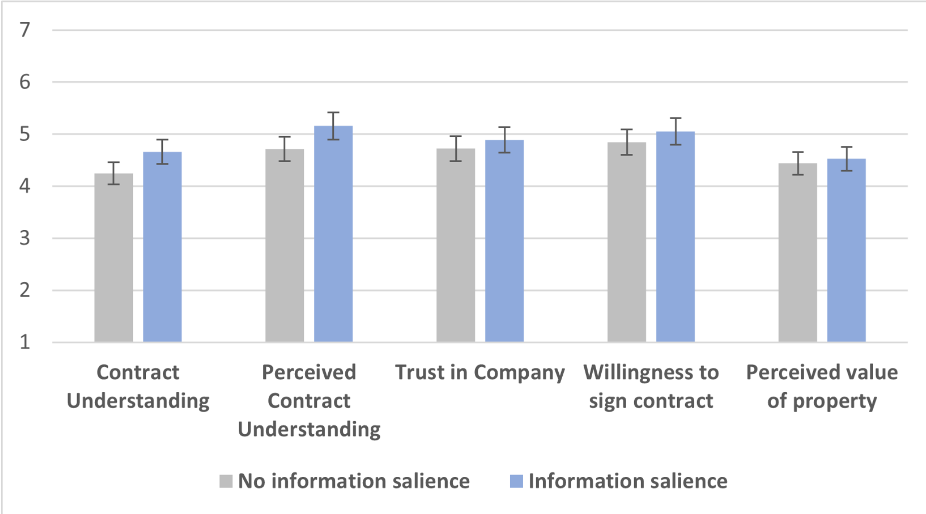

the landlord increased willingness to sign contracts ($p < 0.001$), it did not mediate the effect of information salience on willingness to sign contracts. Hypothesis 3a is not supported.

Hypothesis 4a. Contract risk did not moderate the effect of information salience on willingness to sign contracts ($p = 0.315$). Hence, Hypothesis 4a is not supported.

Figure 3.8 Study 4: The Effect of Contract Risk on Consumer Outcomes

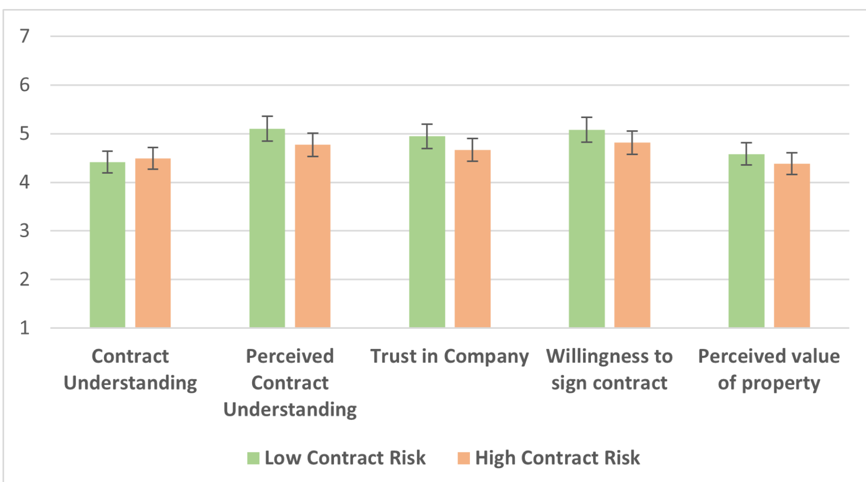

The Effect of Information Salience and Contract Risk on Perceived Value

The full moderated mediation model was significant, $F(15, 438) = 26.914, R^2 = 0.480, p < 0.001$.

Table 3.4 Study 4: Moderated Mediation Results—Willingness to Sign and Perceived Value

	MEDIATORS						OUTCOME			
	Understand.		**P. Understand.**		**Trust**		**Willing. to Sign**		**Perc. Value**	
	coeff	*se*	*coeff*	*se*	*coeff*	*se*	*coeff*	*se*	*coeff*	*se*
Constant	4.16***	0.40	5.82***	0.53	5.12***	0.50	1.02†	0.60	1.97***	0.39
Information salience	0.35*	0.14	0.26	0.18	-0.03	0.17	-0.02	0.17	-0.02	0.11
Understanding							-0.10†	0.06	-0.08**	0.04
Perc. understanding							0.18**	0.06	0.11**	0.04
Trust							0.64***	0.06	0.49***	0.04
Risk	0.02	0.14	-0.35†	0.18	-0.43*	0.17	-0.11	0.17	-0.04	0.11
Interaction (trans x risk)	0.18	0.20	0.15	0.26	0.29	0.25	0.23	0.24	0.06	0.16
Reading habits	0.06	0.04	-0.29***	0.05	-0.13**	0.05	0.06	0.05	0.02	0.03
Gender	0.04	0.10	-0.05	0.13	0.05	0.12	0.02	0.12	-0.07	0.08
Age	0.00	0.01	0.00	0.01	-0.01	0.01	0.00	0.01	0.00	0.00
Language	-0.14	0.24	0.42	0.31	0.45	0.30	0.04	0.29	0.00	0.19
Education	-0.06	0.05	0.00	0.07	-0.08	0.06	0.06	0.06	0.02	0.04
Ethnicity	-0.02	0.03	-0.02	0.04	-0.03	0.04	-0.02	0.04	-0.03	0.03
Relationship	0.00	0.05	0.07	0.07	0.13†	0.07	-0.04	0.06	0.00	0.04
Income	0.01	0.03	0.02	0.04	0.05	0.04	0.03	0.04	0.07	0.03
N	454		454		454		454		454	
R^2	0.060		0.110		0.055		0.424		0.476	
Adj. R^2	0.032		0.084		0.027		0.407		0.460	
Residual Std. Error	1.060		1.882		1.687		1.564		0.676	
F Statistic (df)	2.561*** (11,442)		4.974*** (11,442)		2.340*** (11,442)		23.091*** (14,439)		28.439*** (14,439)	

Note. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; † $p < 0.10$

Hypothesis 2b. As shown above, information salience increased contract understanding ($p < 0.001$) and perceived contract understanding ($p = 0.027$). Both contract understanding and perceived contract understanding mediated the effect of information salience on perceived value ($p = 0.038, p = 0.006$), regardless of trust in the landlord. Hence, Hypothesis 2b is supported.

Hypothesis 3b. As mentioned above, information salience did not significantly increase trust in the landlord ($p = 0.138$). Hence, even though trust in the landlord increased perceived value ($p < 0.001$), it did not mediate the effect of information salience on perceived value. Hypothesis 3b is not supported.

Hypothesis 4b. Contract risk did not moderate the effect of information salience on perceived value ($p = 0.747$). Hypothesis 4b is not supported.

Hypothesis 4c. Contract risk did not moderate the effect of information salience on contract understanding ($p = 0.350$), nor on perceived contract understanding ($p = 0.556$). Hypothesis 4c is not supported.

Hypothesis 4d. Contract risk did not moderate the effect of information salience on trust in the company ($p = 0.234$). Hypothesis 4d is not supported.

7.4. Discussion

Study 4 explores the effect of information salience in an apartment leasing context, which could potentially pose higher, more personal risks. Results in Study 4 again suggest that information salience in contracts can increase positive perceptions of consumers. We show that perceived contract understanding explains the positive effect of information salience on consumers' willingness to sign contracts and perceived value of the property, consistent with Study 3. However, this time, actual contract understanding mediated the effect of information salience on consumer outcomes. Nevertheless, trust (in the landlord) did not explain this effect unlike Study 3. Results again show the attenuating role of high contract risk in these positive effects. However, for the leasing context, we could not find significant moderating effect of contract risk.

8. General Discussion

8.1. Summary of Results

Our studies show that providing information salience into contracts by displaying a summary of each paragraph can be an effective means to increase consumer positive perceptions of contract and company. We observe that consumer contract understanding and trust in the company explain the positive effect of information salience on consumer perceptions of contract and company. Interestingly, in the leasing context (which could pose higher risks), showing the summaries did not increase consumer trust. Coupled with this finding, we find that consumers' responses to the information salience varies with the level of risk associated with the contract, e.g., whether user's privacy is protected in social media.

8.2. Managerial Implications

Our results illustrate that information salience in contracts—or highlighting consumer contracts—can benefit both consumers and firms. We leave additional remarks that could help guide the managers to show contract summaries in a way that benefits both firms and consumers.

First, it is important for managers to display the summaries together with the original version of consumer contracts. In many jurisdictions around the world, it may simply not be legally feasible to disclose just the nutshell version of a consumer contract, as doing so may result in leaving important legal nuances present in the original version. In addition, given that vast majority of consumers are deeply acquainted with the current long, convoluted version of the consumer contracts (Furth-Matzkin 2018), leaving out the original form of consumer contracts may leave consumers with questions about whether viewing just the translated version carries additional legal risks and implications. In addition, by presenting both versions, consumers can always refer back to the original version of the contract whenever they want to learn more about the details of particular clause. Information salience is designed to help pique consumers' initial interest, but it is not designed to encapsulate information in its entirety or substitute the original version of the contract. Disclosing the summaries alongside the original version, thus, can benefit both firms and customers alike.

Second, actual content of the contract matters. Our findings show that trust is necessary for information salience for firms to signal trustworthiness and to form positive consumer perceptions. If consumer contracts were, in fact, harmful or unfair to the consumers, showing the summaries alongside the original version of the contract may actually backfire. Moreover, companies may want to deliberate carefully about the overall tone of the key points highlighted in the summary. Although it remains in our future study to compare the effect of casual/warm vs. formal/rigid language, we believe that the tone of the summary should be designed to bring the companies and consumers psychologically closer together and provide consumers opportunities to develop greater trust toward the firms.

Third, success of summary disclosure hinges on the degree of consumer agency and empowerment. If the summaries only list out restrictions or companies' liability limitations without laying out what rights customers have, customers may not have positive perceptions toward the company or the contract. To the greatest extent possible, consumer contract and its summaries should be written in a way that empowers consumers, not belittle them.

8.3. Limitations and Future Research

There are many future directions to which our research can be extended. A natural extension would be to implement our study in a field setting. One limitation of testing in a laboratory setting is that, since the subjects in the lab anticipate answering questions about consumer contracts, they are more likely to pay attention to it than they usually would. Complementing our studies with the field study would provide external validity and credence of positive effects of information salience.

Similarly, the research can be extended to study how information salience affects consumers behaviorally, especially when the contracts contain terms that directly, imminently impacts consumers' well-being. For instance, consumers who invest in cryptocurrency might not initially read the terms and conditions at the time of purchase, but might choose to read it once the consumer becomes aware of the extremely volatile risk of trading cryptocurrency.

One could also study how different lingual syntax and tones of the contract could affect consumers' perception of consumer contract and the firm. For instance, the study could examine whether using casual tone narrows the psychological distance between consumers and firms. On the other hand, when the contract is written in a casual language but is overly preferential to the company, consumers' trust may dissipate more quickly.

References for Chapter 1

Anderson E (1994) Cross-Category Variation in Customer Satisfaction and Retention. *Mark. Lett.* 5(1):19–30.

Anderson ET, Simester DI (2004) Long-run effects of promotion depth on new versus established customers: Three field studies. *Mark. Sci.* 23(1):4–20.

Anderson EW, Fornell C, Lehmann DR (1994) Customer Satisfaction, Market Share, and Profitability: Findings from Sweden. *J. Mark.* 58(3):53–66.

Anderson RE, Jolson MA (1980) Technical wording in advertising: Implications for market segmentation. *J. Mark.* 44(Winter):57–66.

Angrist JD, Pischke JS (2008) *Mostly harmless econometrics: An empiricist's companion* (Princeton University Press).

Ayres I, Schwartz A (2014) The no-reading problem in consumer contract law. *Stanford Law Rev.* 66(3):545–610.

Bakos Y, Marotta-Wurgler F, Trossen DR (2014) Does anyone read the fine print? Consumer attention to standard-form contracts. *J. Legal Stud.* 43(1):1–35.

Blake T, Moshary S, Sweeney K, Tadelis S (2018) Price salience and product choice. Working Paper, National Bureau of Economic Research, Boston.

Bollinger B, Leslie P, Sorensen A (2011) Calorie posting in chain restaurants. *Am. Econ. J. Econ. Policy* 3(1):91–128.

Buell RW, Campbell D, Frei FX (2016) How do customers respond to increased service quality competition? *Manuf. Serv. Oper. Manag.* 18(4):1–33.

Buell RW, Campbell D, Frei FX (2020) The customer may not always be right: Customer compatibility and service performance. *Manage. Sci.* (Forthcoming).

Buell RW, John LK (2018) Commonwealth Bank of Australia: Unbanklike Experimentation (619-018). Harvard Business School case 619-018, Harvard Business School Publishing, Boston.

Buell RW, Kalkanci B (2020) How transparency into internal and external responsibility initiatives influences consumer choice. *Manage. Sci.* (Forthcoming).

Buell RW, Kim T, Tsay CJ (2017) Creating reciprocal value through operational transparency. *Manage. Sci.* 63(6):1673–1695.

Buell RW, Norton MI (2011) The labor illusion: How operational transparency increases perceived value. *Manage. Sci.* 57(9):1564–1579.

Buell RW, Porter E, Norton MI (2020) Surfacing the submerged state: Operational transparency increases trust in and engagement with government. *Manuf. Serv. Oper. Manag.* (Forthcoming).

Buell RW, Campbell D, Frei FX (2010) Are self-service customers satisfied or stuck? *Prod. Oper. Manag.* 12(6):679–697.

Cabral LL, Hortacsu A, Hortaçsu A (2010) The dynamics of seller reputation: Evidence from eBay. *J. Ind. Econ.* 58(1):54–78.

Cachon GP, Harker PT (2002) Competition and outsourcing with scale economies. *Manage. Sci.* 48(10):1314–1333.

Chetty R, Looney A, Kroft K (2009) Salience and taxation: Theory and evidence. *Am. Econ. Rev.* 99(4):1145–1177.

Cohen MA, Whang S (1997) Competing in product and service: A product life-cycle model. *Manage. Sci.* 43(4):535–545.

Coyles S, Gokey T (2005) Customer retention is not enough. *J. Consum. Mark.* 22(2):101–105.

Derlega VJ (1984) Self-disclosure and intimate relationships. *Commun. intimacy, close relationships*. (Academic Press), 230.

Finkelstein A (2009) E-ztax: Tax salience and tax rates. *Q. J. Econ.* 124(3):969–1010.

Frei FX (2006) Breaking the trade-off between efficiency and service. *Harv. Bus. Rev.* 84(11):93–101.

Gabszewicz J, Thisse JFJF (1979) Price competition, quality and income disparities. *J. Econ. Theory* 359(3):340–359.

Gans N (2002) Customer loyalty and supplier quality competition. *Manage. Sci.* 48(2):207.

Gourville JT, Soman D (2005) Overchoice and assortment type: When and why variety backfires. *Mark. Sci.* 24(3):382–395.

Guajardo JA, Cohen MA (2017) Service differentiation and operating segments: A framework and an application to after-sales services. *Manuf. Serv. Oper. Manag.* Forthcomin:38.

Hall J, Porteus E (2000) Customer service competition in capacitated systems. *Manuf. Serv. Oper. Manag.* 2(2):144.

Heskett J, Sasser E, Schlesinger L (1997) *The Service Profit Chain* (The Free Press, New York).

Ho TH, Zheng YS (2004) Setting customer expectation in service delivery: An integrated marketing-operations perspective. *Manage. Sci.* 50(4):479–488.

Hoffman-Graff MA (1977) Interviewer use of positive and negative self-disclosure and interviewer-subject sex pairing. *J. Couns. Psychol.* 24(3):184–190.

Israel M (2005) Services as experience goods : An empirical examination of consumer learning in automobile insurance. *Am. Econ. Rev.* 95(5):1444–1463.

Iyengar SS, Lepper MR (2000) When choice is demotivating: Can one desire too much of a good thing? *J. Pers. Soc. Psychol.* 79(6):995–1006.

John LK, Barasz K, Norton MI (2016) Hiding personal information reveals the worst. *Proc. Natl. Acad. Sci.* 113(4):954–959.

Johnson EJ, Russo JE (1984) Product familiarity and learning new information. *J. Consum. Res.* 11(1):542–550.

Jones T, Sasser E (1995) Why satisfied customers defect. *Harv. Bus. Rev.* 73(6):88–99.

Jourard SM (1971) *Self-disclosure: An experimental analysis of the transparent self* (John Wiley & Sons, Oxford, England).

Kalkanci B, Ang E, Plambeck E (2018) Managing social and environmental impacts with voluntary vs. mandatory disclosure to investors. *Manage. Sci.* (Forthcoming).

Karmarkar U, Pitbladdo R (1995) Service markets and competition. *J. Oper. Manag.* 12(3):397–411.

Karmarkar US, Pitbladdo RC (1997) Quality, class, and competition. *Manage. Sci.* 43(1):27–39.

Kivetz R, Simonson I (2000) The effects of incomplete information on consumer choice. *J. Mark. Res.* 37(4):427–448.

Kosack S, Fung A (2014) Does transparency improve governance? *Annu. Rev. Polit. Sci.* 17:65–87.

Krishnamurthi L, Raj SP (2006) The effect of advertising on consumer price sensitivity. *J. Mark. Res.* 22(1985):119–129.

LaBarbera P, Mazursky D (1983) A longitudinal assessment of consumer satisfaction / dissatisfaction: The dynamic aspect of the cognitive process. *J. Mark. Res.* 20(4):393–404.

Laurenceau JP, Barrett LF, Pietrornonaco PR (1998) Intimacy as an interpersonal process: The importance of self-disclosure, partner disclosure, and perceived partner responsiveness in interpersonal exchanges. *J. Pers. Soc. Psychol.* 74(5):1238–1251.

Lewis M (2006) Customer acquisition promotions and customer asset value. *J. Mark. Res.* 43(2):195-203.

Li L, Lee YS (1994) Pricing and delivery-time performance in a competitive environment. *Manage. Sci.* 40(5):633–646.

Luca M (2016) Reviews, reputation, and revenue: The case of Yelp.com. Working Paper 12-016, Harvard Business School, Boston.

Luca M, Smith J (2013) Salience in quality disclosure: Evidence from the U.S. News college rankings. *J. Econ. Manag. Strateg.* 22(1):58–77.

McDougall GHG, Levesque T (2000) Customer satisfaction with services: Putting perceived value into the equation. *J. Serv. Mark.* 14(5):392–410.

Mohan B, Buell RW, John LK (2020) Lifting the veil: The benefits of cost transparency. *Mark. Sci.* 39(6):1033–1201.

Moon Y (2000) *Intimate exchanges: Using computers to elicit self-disclosure from consumers*

Obar JA, Oeldorf-Hirsch A (2020) The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Inf. Commun. Soc.* 23(1):128–147.

Oliver RL (1993) Cognitive, affective, and attribute bases of the satisfaction response. *J. Consum. Res.* 20(3):418.

Parasuraman A, Zeithaml VA, Berry LL (1985) A conceptual model of service quality and its implications for future research. *J. Mark.* 49(4):41–50.

Perlman Daniel, Fehr B (1987) The development of intimate relationships. Perlman D, Duck S, eds. *Intim. relationships Dev. Dyn. Deterior.* (Sage Publications, Inc.), 13–42.

Reichheld F, Sasser E (1990) Zero defections: Quality comes to services. *Harv. Bus. Rev.* 68(5):105–111.

Roels G (2014) Optimal design of coproductive services: Interaction and work allocation. *Manuf. Serv. Oper. Manag.* 16(4):578–594.

Sampson SE, Froehle CM (2006) Foundations and implications of a proposed unified services theory. *Prod. Oper. Manag.* 15(2):329–343.

Schmidt Wi, Buell RW (2017) Experimental evidence of pooling outcomes under information asymmetry. *Manage. Sci.* 63(5):1586–1605.

Shaked A, Sutton J (1982) Relaxing price competition through product differentiation. *Econ. Anal.* 49(1):3–13.

Sutton J (1986) Vertical product differentiation: Some basic themes. *Am. Econ. Rev.* 76(2):393–398.

Tansik DA, Chase RBR, Tansik DA (1983) The customer contact model for organization design. *Manage. Sci.* 29:1037–1051.

Tirole J (1990) *The theory of industrial organization* Fourth. (The MIT Press, Cambridge, MA).

Tsay AA, Agarwal N (2000) Channel dynamics under price and service competition. *Manuf. Serv. Oper. Manag.* 2(4):372–391.

Tse D, Wilton P (1988) Models of consumer satisfaction formation: an extension. *J. Mark.* 25(2):204-212.

Wheeless LR, Grotz J (1977) The measurement of trust and its relationship to self disclosure. *Hum. Commun. Res.* 3(3):250–257.

Xue M, Harker P (2002) Customer efficiency: Concept and its impact on e-business management. *J. Serv. Res. JSR* 4(4):253–268.

Yang Z, Peterson RT (2004) Customer perceived value, satisfaction and loyalty: The role of switching costs. *Psychol. Mark.* 21(10):799–822.

Disclosure for Chapter 1

Funding for this research, apart from the costs Commonwealth Bank of Australia internally incurred to support the experiment, was provided by Harvard Business School. Ryan Buell has received compensation from Commonwealth Bank of Australia in the past for executive education teaching.

Acknowledgements for Chapter 1

We gratefully acknowledge and thank the Behavioural Economics team at Commonwealth Bank of Australia for their partnership in designing and executing this field experiment, including Rafael Batista, Nathan Damaj, Emily Daniels, Mohamed Khalil, Will Mailer, Jessica Min, and Nikhil Ravichandar. We are also grateful to the many efforts of people throughout the organization whose time, consideration, and care were essential to the project, including Elisa Barnard, Jenni Barnett, Grant Burrow, Ian Calbert, Keith Callus, Tom Carr, Valentina Ciampi, Lauren Cox, Kate Crous, Amy Cunningham, Tina DeLorenzo, Belinda Duff, Jennifer Excell, Zomi Frankcom, Beau Fraser, Ben Freeman, Sheffy Goyal, Ben Grauer, Chris Hill, Nilesh Karunaratne, Emir Kazazic, Diane Loo, Lydia MacDonald, Rebecca Manly, Fiona Prater, Katie Roch, David Singh, Mary Yang and Jade Young. We're also grateful for the support and guidance of Alain Bonacossa, Jean Cunningham, Chip Flowers, Robert Freeman, Catherine Nehring, and Christopher Pringle, at Harvard University, who helped to make this research collaboration possible, as well as faculty affiliates in the Harvard Sustainability Transparency and Accountability Research (STAR) Lab, including John Beshears, Rebecca Henderson, Leslie John, David Laibson, Brigitte Madrian, Sendhil Mullainathan, and Michael Norton, who supported this work in various ways. We are especially grateful to Michael

Hiscox, who directs the STAR Lab and brought us all together in the first place. Finally, we thank Sophie Bick, Kim Jacobson, and Elizabeth Leh, for their support, and Melissa Ouellet, whose many contributions were instrumental in our analysis. This research also benefited tremendously from seminar participants at Cornell University, Harvard Business School, Johns Hopkins University, Massachusetts Institute of Technology, and Purdue University, as well as participants in the Harvard Business School TOM Unit Doctoral Seminar, the 2018 Winter Operations Conference, the 2018 MSOM Annual Conference, the 2018-2020 INFORMS Annual Conferences, and the 2018-2020 POMS Annual Conferences, whose constructive questions and suggestions influenced its development. All errors remain our own.

References for Chapter 2

Angrist JD, Pischke JS (2008) *Mostly harmless econometrics: An empiricist's companion* (Princeton university press).

Bray R (2019) Operational transparency: Showing when work gets done. *Working Paper.*

Breiman L (2001) Random forests. *Machine learning* 45(1):5–32.

Breiman L, Friedman J, Olshen R, Stone C (1984) Classification and regression trees. Wadsworth Int. *Group* 37(15):237–251.

Caro F, Gallien J (2012) Clearance pricing optimization for a fast-fashion retailer. *Operations Research* 60(6):1404–1422.

Cui R, Gallino S, Moreno A, Zhang DJ (2018) The operational value of social media information. *Production and Operations Management* 27(10):1749–1769.

Cui R, Li M, Li Q (2019) Value of high-quality logistics: Evidence from a clash between sf express and alibaba. *Management Science* 66(9):3799–4358.

Cui R, Sun T, Lu Z, Golden J (2020) Sooner or later? promising delivery speed in online retail. *Working Paper.*

Ferreira KJ, Lee BHA, Simchi-Levi D (2016) Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18(1):69–88.

Fisher ML, Gallino S, Xu JJ (2019) The value of rapid delivery in omnichannel retailing. *Journal of Marketing Research* 56(5):732–748.

Guo X, Grushka-Cockayne Y, De Reyck B (2019) Forecasting airport transfer passenger flow using real-time data and machine learning. *Working Paper.*

Gutierrez GJ, Kouvelis P (1991) Parkinson's law and its implications for project management. *Management Science* 37(8):990–1001.

Hall RW (1991) *Queueing methods: for services and manufacturing* (Pearson College Division).

Hasija S, Pinker EJ, Shumsky RA (2008) Call center outsourcing contracts under information asymmetry. *Management Science* 54(4):793–807.

Ibanez MR, Clark JR, Huckman RS, Staats BR (2018) Discretionary task ordering: Queue management in radiological services. *Management Science* 64(9):4389–4407.

Imbens GW, Lemieux T (2008) Regression discontinuity designs: A guide to practice. *Journal of econometrics* 142(2):615–635.

James G, Witten D, Hastie T, Tibshirani R (2013) *An introduction to statistical learning*, volume 112 (Springer).

JD.com (2020) jd.com - authentic products delivered today. https://corporate.jd.com/ourBusiness, accessed: 2020-09-30.

Kc DS, Terwiesch C (2012) An econometric analysis of patient flows in the cardiac intensive care unit. *Manufacturing & Service Operations Management* 14(1):50–65.

Kim Y, Yano CA (1994) Minimizing mean tardiness and earliness in single-machine scheduling problems with unequal due dates. *Naval Research Logistics* 41(7):913–933.

Lee DS, Lemieux T (2010) Regression discontinuity designs in economics. *Journal of economic literature* 48(2):281–355.

Lu Y, Musalem A, Olivares M, Schilkrut A (2013) Measuring the effect of queues on customer purchases. *Management Science* 59(8):1743–1763.

Mandelbaum A, Shimkin N (2000) A model for rational abandonments from invisible queues. *Queueing Systems* 36(1-3):141–173.

Mao W, Ming L, Rong Y, Tang CS, Zheng H (2019) Faster deliveries and smarter order assignments for an on-demand meal delivery platform. *Working Paper*.

Marin CV, Drury CG, Batta R, Lin L (2007) Human factors contributes to queuing theory: Parkinson's law and security screening. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 51, 602–606 (SAGE Publications Sage CA: Los Angeles, CA).

Musalem A, Shang Y, Song JSJ (2017) An empirical study of customer spillover learning about service quality. *Working Paper*.

Newell C (2013) *Applications of queueing theory*, volume 4 (Springer Science & Business Media).

Parkinson CN, Osborn RC (1957) *Parkinson's law, and other studies in administration*, volume 24 (Houghton Mifflin Boston).

Perakis G, Sim M, Tang Q, Xiong P (2018) Joint pricing and production: a fusion of machine learning and robust optimization. *Working Paper*.

Pinedo M (2012) *Scheduling*, volume 29 (Springer).

Raileanu LE, Stoffel K (2004) Theoretical comparison between the gini index and information gain criteria. *Annals of Mathematics and Artificial Intelligence* 41(1):77–93.

Rao S, Rabinovich E, Raju D (2014) The role of physical distribution services as determinants of product returns in internet retailing. *Journal of Operations Management* 32(6):295–312.

Shen ZJM, Tang CS, Wu D, Yuan R, Zhou W (2020) JD. com: Transaction level data for the 2020 MSOM data driven research challenge. *Forthcoming, Manufacturing & Service Operations Management*.

Varey C, Kahneman D (1992) Experiences extended across time: Evaluation of moments and episodes. *Journal of Behavioral Decision Making* 5(3):169–185.

Yu Q, Zhang Y, Zhou YP (2020) Delay information in virtual queues: A large-scale field experiment on a ride-sharing platform. *Working Paper*.

Acknowledgments for Chapter 2

The authors of this manuscript are grateful to JD.com and the M&SOM Society for making the datasets available. The authors would like to express much gratitude to Antonio Moreno for providing invaluable feedback and detailed comments throughout the development of this paper. The paper also benefited immensely from comments and thoughts from various seminar participants.

References for Chapter 3

Acquisti A, John LK, Loewenstein G (2013) What is privacy worth? *Journal of Legal Studies* 42(2):249–274, URL http://dx.doi.org/10.1086/671754.

Aguirre E, Mahr D, Grewal D, de Ruyter K, Wetzels M (2015) Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing* 91(1):34–49, URL http://dx.doi.org/10.1016/j.jretai.2014.09.005.

Atkinson L, Rosenthal S (2014) Signaling the green sell: The influence of eco-label source, argument specificity, and product involvement on consumer trust. *Journal of Advertising* 43(1):33–45, URL http://dx.doi.org/10.1080/00913367.2013.834803.

Ayres I, Schwartz A (2014) The no-reading problem in consumer contract law. *Stanford Law Review* 66(3):545–610, ISSN 00389765.

Bakos Y, Marotta-Wurgler F, Trossen DR (2014) Does anyone read the fine print? Consumer attention to standard-form contracts. *Journal of Legal Studies* 43(1):1–35, URL http://dx.doi.org/10.1086/674424.

Bebchuk LA, Posner RA (2005) One-sided contracts in competitive consumer markets. *Michigan Law Review* 104:824.

Becher SI, Zarsky TZ (2008) E-contract doctrine 2.0: standard form contracting in the age of online user participation. *Michigan Telecommunications and Technology Law Review* 14:303–366.

Ben-Shahar O, Schneider CE (2011) The failure of mandated disclosure. *University of Pennsylvania Law. Review* 159(3):647–749, URL http://dx.doi.org/10.23943/princeton/9780691161709.003.0003.

Berreby D (2017) Click to agree with what? No one reads terms of service, studies confirm. URL: https://www.theguardian.com/technology/2017/mar/03/terms-of-service-online-contracts-fine-print.

Bertrand M, Karlan D, Mullainathan S, Shafir E, Zinman J (2010) What's advertising content worth? Evi- dence from a consumer credit marketing field experiment. *Quarterly Journal of Economics* 125(1):263– 306.

Best A (1981) *When consumers complain* (New York), URL http://dx.doi.org/10.1177/0092070396244006.

Blake T, Moshary S, Sweeney K, Tadelis S (2018) Price Salience and Product Choice. *Marketing Science* Articles i:1–18.

Bollinger B, Leslie P, Sorensen A (2011) Calorie posting in chain restaurants. *American Economic Journal: Economic Policy* 3(1):91–128.

Bosker B (2010) 7,500 Online Shoppers Accidentally Sold Their Souls To Gamestation. URL http://dx.doi.org/10.1503/cmaj.071733.

Boulding W, Kirmani A (1993) A Consumer-Side Experimental Examination of Signaling Theory: Do Consumers Perceive Warranties as Signals of Quality? *Journal of Consumer Research* 20(1):111, URL http://dx.doi.org/10.1086/209337.

Brough AR, Norton DA, John L (2020) The Bulletproof Glass Effect: When Privacy Notices Backfire. Working Paper.

Buell RW, Choi M (2021) Improving Customer Compatibility with Operational Transparency. Working Paper.

Buell RW, Kim T, Tsay CJ (2017) Creating Reciprocal Value Through Operational Transparency. *Management Science* 63(6):1673–1695.

Buell RW, Norton MI (2011) The labor illusion: How operational transparency increases perceived value. *Management Science* 57(9):1564–1579.

Buell RW, Porter E, Norton MI (2020) Surfacing the Submerged State: Operational Transparency Increases Trust in and Engagement with Government. *Manufacturing & Service Operations Management*.

Chetty R, Looney A, Kroft K, Akerlof G, Ahn D, Auerbach A, Carpenter K, Della Vigna S, Finkelstein A, Greenstone M, Kariv S, Katuscak P, Koszegi B, Luttmer E, Poterba J, Rabin M, Reis R, Saez E, Shapiro J, Shleifer A, Simonsohn U (2009) Salience and Taxation: Theory and Evidence. *American economic review* 99(4):1145–77.

Connelly BL, Certo ST, Ireland RD, Reutzel CR (2011) Signaling theory: A review and assessment. *Journal of Management* 37(1):39–67, URL http://dx.doi.org/10.1177/0149206310388419.

Debo L (2010) Prices and Congestion as Signals of Quality. Technical report.

Derlega VJ (1984) Self-disclosure and intimate relationships. *Communication, intimacy, and close relationships* 1–9.

Dickens RL (2007) Finding Common Ground in the World of Electronic Contracts: The Consistency of Legal Reasoning in Clickwrap Cases. *Marq. Intell. Prop. L. Rev* 11:379.

Donnelly GE, Zatz LY, Svirsky D, John LK (2018) The Effect of Graphic Warnings on Sugary-Drink Purchasing. *Psychological Science* 29(8):1321–1333, URL http://dx.doi.org/10.1177/0956797618766361.

Eisenberg MA (1995) The Limits of Cognition and the Limits of Contract. *Stanford Law Review* 47(2):211, URL http://dx.doi.org/10.2307/1229226.

Eisingerich AB, Bell SJ (2008) Perceived service quality and customer trust: Does enhancing customers' service knowledge matter? *Journal of Service Research* 10(3):256–268, URL http://dx.doi.org/10.1177/1094670507310769.

Furth-Matzkin M (2017) On the unexpected use of unenforceable contract terms: Evidence from the residential rental market. *Journal of Legal Analysis* 9(1):1–49, URL http://dx.doi.org/10.1093/jla/ lax002.

Furth-Matzkin M (2018) the Harmful Effects of Unenforceable Contract Terms: Experimental Evidence. *Ala. Law Review* 70:1031–1071.

Furth-Matzkin M, Sommers R (2020) Consumer psychology and the problem of fine-print fraud. *Stanford Law Review* 72(3):503–560.

Goodwin C (1991) Privacy: Recognition of a Consumer Right. *Journal of Public Policy & Marketing* 10(1):149–166, URL http://dx.doi.org/10.1177/074391569101000111.

Groth M, Gilliland SW (2006) Having to wait for service: Customer reactions to delays in service delivery. *Applied Psychology* 55(1):107–129, URL http://dx.doi.org/10.1111/j.1464-0597.2006.00238.x.

Hayes AF (2017) *Introduction to Mediation, Moderation, and Conditional Process Analysis*.

Hillman RA (2006) Online Boilerplate: Would Mandatory Website Disclosure of E-Standard Terms Backfire? *Michigan Law Review* 104(5).

Hoffman DL, Novak TP, Peralta M (1999) Building Consumer Trust Online. *Communications of the ACM* 42(4):80–85, URL http://dx.doi.org/10.1145/299157.299175.

John LK, Acquisti A, Loewenstein G (2011) Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *Journal of Consumer Research* 37(5):858–873, URL http://dx.doi.org/10.1086/656423.

Kim K, Kim J (2011) Third-party Privacy Certification as an Online Advertising Strategy: An Investigation of the Factors Affecting the Relationship between Third-party Certification and Initial Trust. *Journal of Interactive Marketing* 25(3):145–158, URL http://dx.doi.org/10.1016/j.intmar.2010.09.003.

Korobkin R (2003) Bounded Rationality, Standard Form Contracts, and Unconscionability. Technical Report 4.

Lacko JM, Pappalardo JK (2004) The Effect of Mortgage Broker Compensation Disclosures on

Consumers and Competition: A Controlled Experiment. Technical Report February.

Laurenceau JP, Feldman Barrett L, Pietromonaco PR (1998) Intimacy as an Interpersonal Process" The Im- portance of Self-Disclosure, Partner Disclosure, and Perceived Partner Responsiveness in Interpersonal Exchanges. Technical Report 5.

Mann RJ, Siebeneicher T (2008) Just one click: The reality of internet retail contracting. *Columbia Law Review* 108(4):984–1012, URL http://dx.doi.org/10.2139/ssrn.988788.

Marotta-Wurgler F (2011) Will increased disclosure help? Evaluating the recommendations of the ALI's "principles of the law of software contracts". *University of Chicago Law Review* 78(1):165–186.

Martin (2016) Understanding Privacy Online: Development of a Social Contract Approach to Privacy. *Journal of Business Ethics* 137(3):551–569, URL http://dx.doi.org/10.1007/s10551-015-2565-9.

Martin (2018) The penalty for privacy violations: How privacy violations impact trust online. *Journal of Business Research* 82(2018):103–116, URL http://dx.doi.org/10.1016/j.jbusres.2017.08.034.

Martin KD, Murphy PE (2017) The role of data privacy in marketing. *Journal of the Academy of Marketing Science* 45(2):135–155, URL http://dx.doi.org/10.1007/s11747-016-0495-4.

Milne GR, Bahl S, Rohm A (2008) Toward a framework for assessing covert marketing practices. *Journal of Public Policy and Marketing* 27(1):57–62, URL http://dx.doi.org/10.1509/jppm.27.1.57.

Milne GR, Culnan MJ (2004) Strategies for reducing online privacy risks: Why consumers read (or don't read) online privacy notices. *Journal of Interactive Marketing* 18(3):15–29, URL http://dx.doi.org/ 10.1002/dir.20009.

Miyazaki AD, Krishnamurthy S (2002) Internet seals of approval: Effects on online privacy policies and consumer perceptions. *Journal of Consumer Affairs* 36(1):28–49, URL http://dx.doi.org/10.1111/ j.1745-6606.2002.tb00419.x.

Mohan B, Buell RW, John LK (2020) Lifting the Veil: The Benefits of Cost Transparency. *Marketing Science* 39(6):1105–1121, URL http://dx.doi.org/10.2139/ssrn.2498174.

Noar SM, Francis DB, Bridges C, Sontag JM, Brewer NT, Ribisl KM (2017) Effects of Strengthening Cigarette Pack Warnings on Attention and Message Processing: A Systematic Review. *Journalism and Mass Communication Quarterly* 94(2):416–442, URL http://dx.doi.org/10.1177/1077699016674188.

Obar JA, Oeldorf-Hirsch A (2018) The Clickwrap: A Political Economic Mechanism for Manufacturing Consent on Social Media. *Social Media and Society* 4(3), ISSN 20563051, URL http://dx.doi.org/ 10.1177/2056305118784770.

Obar JA, Oeldorf-Hirsch A (2019) The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* 1–20.

Olivero N, Lunt P (2004) Privacy versus willingness to disclose in e-commerce exchanges: The effect of risk awareness on the relative role of trust and control. *Journal of Economic Psychology* 25(2):243–262, URL http://dx.doi.org/10.1016/S0167-4870(02)00172-1.

Rust RT, Verhoef PC (2005) Optimizing the marketing interventions mix in intermediate-term CRM. *Marketing Science* 24(3):477–489, URL http://dx.doi.org/10.1287/mksc.1040.0107.

Schmidt W, Buell RW (2017) Experimental evidence of pooling outcomes under information asymmetry. *Management Science* 63(5):1586–1605, ISSN 15265501, URL http://dx.doi.org/10.1287/mnsc.2015.2407.

Schwartz A, Wilde LL (1983) Imperfect information in markets for contract terms: The examples of warranties and security interests. *Virginia Law Review* 1387–1485.

Shu SB (2010) Consumer Behavior in Relation to Truth in Lending Disclosures. *First Annual Boulder Summer Conference on Consumer Financial Decision Making*.

Silverman C, Wilson JL (2016) State Attorney General Enforcement of Unfair or Deceptive Acts and Practices Laws: Emerging Concerns and Solutions. *Kansas Law Review* 65:209–269, URL http://dx.doi.org/10.17161/1808.25552.

Sovern J, Greenberg EE, Kirgis PF, Liu Y (2014) 'Whimsy Little Contracts' with Unexpected Consequences: An Empirical Analysis of Consumer Understanding of Arbitration Agreements. *SSRN Electronic Journal* 75(1), URL http://dx.doi.org/10.2139/ssrn.2516432.

Spence M (1973) Job Market Signaling. *The Quarterly Journal of Economics* 87(3):355–374.

Spence M (2002) Signaling in Retrospect and the Informational Structure of Markets. *The American Economic Review* 92(3):434–459.

Stark DP, Choplin J (2009) A License to Deceive: Enforcing Contractual Myths Despite Consumer Psycho- logical Realities. *NYU J. L. & Bus.* 5:617.

Steinfeld N (2016) "I agree to the terms and conditions": (How) do users read privacy policies online? An eye-tracking experiment. *Computers in Human Behavior* 55:992–1000, URL http://dx.doi.org/10.1016/j.chb.2015.09.038.

Stewart KJ (2003) Trust transfer on the World Wide Web. *Organization Science* 14(1): 5–17.

Sweeney J, Soutar G (2001) Customer perceived value: The development of a multiple item scale in hospitals. *Journal of Retailing* 203–220.

Sweeny K, Melnyk D, Miller W, Shepperd JA (2010) Information Avoidance: Who, What, When, and Why. *Review of General Psychology* 14(4):340–353, URL http://dx.doi.org/10.1037/a0021288.

Trifts V, Häubl G (2003) Information availability and consumer preference: Can online retailers benefit from providing access to competitor price information? *Journal of Consumer Psychology* 13(1-2):149–159.

Urban GL (2004) The emerging era of customer advocacy. *MIT Sloan Management Review* 45(2):26–32.

Volokh A (2008) Choosing interpretive methods: A positive theory of judges and everyone else. *New York University Law Review* 83(3):769–846, ISSN 00287881.

White AM, Mansfield CL (2002) Literacy and contract. *Stanford Law and Policy Review* 13(233).

Wilkinson-Ryan T (2014) A psychological account of consent to fine print. *Iowa Law Review* 99(4):1745–1784, ISSN 00210552.

Wu KW, Huang SY, Yen DC, Popova I (2012) The effect of online privacy policy on consumer privacy concern and trust. *Computers in Human Behavior* 28(3):889–897, URL http://dx.doi.org/10.1016/ j.chb.2011.12.008

Acknowledgments for Chapter 3

We would like to express many thanks to Antonio Moreno, Leslie John, and past participants of seminars and conferences for constructive and invaluable feedback. We would also like to thank Harvard Business School Division of Research and Faculty Development for providing funding to support our research.

Online appendix for "Improving Customer Compatibility with Tradeoff Transparency"

Visualizing Prior Category Experience

We conduct an analysis of all customers visiting the CBA credit card marketing website, irrespective of assignment to experimental condition, for which information about whether they had a credit card was known. As the information on whether customers had credit cards was collected after their visit, and thus may have included cards opened during the period of the experiment, we do not use this variable as a control in our analysis. However, it enables us to model how the probability of having a credit card varies with age. We use the following logistic regression specification, where AGE_i is entered into the specification as a categorical variable.

$$\Pr\left(CARD_i\right) = f\left(\alpha_0 + \alpha_1 AGE_i + \epsilon_i\right) \quad (1)$$

Next, we plot the marginal effects for each age category, calculating the 95% confidence interval band, as shown in Figure A1.1. The results indicate that the probability of having a credit card increases non-linearly in customer age.

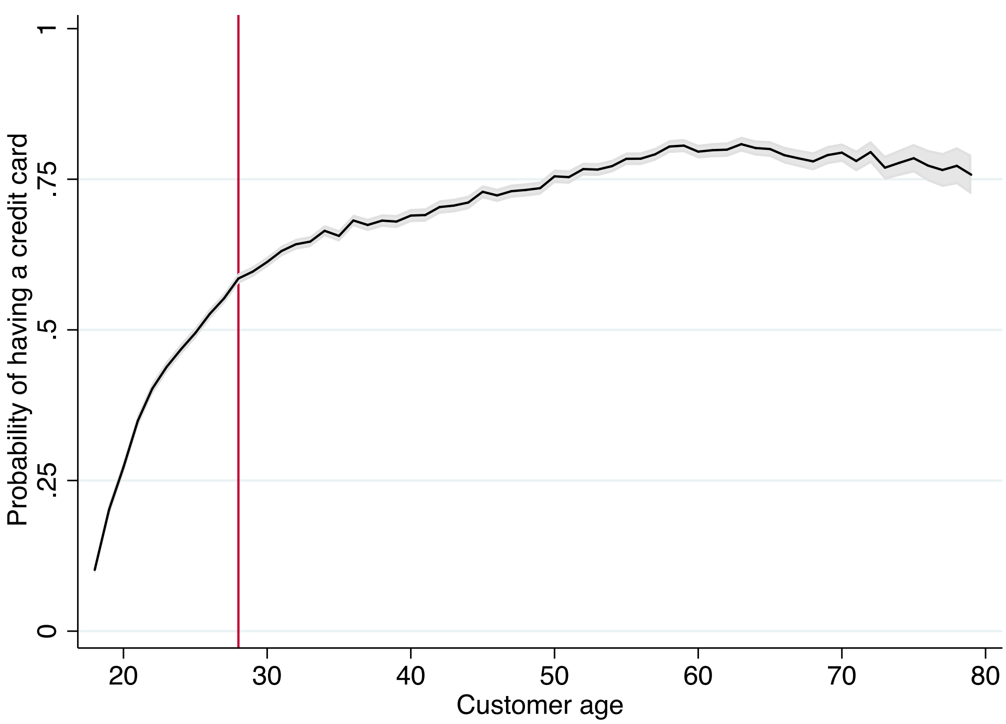

Figure A1.1: The probability a prospective customer shopping for a credit card already has a credit card, as a function of their age (n=389,085). Predicted plots are based on a logistic regression of the probability of having a credit card as a function of indicator variables for the prospective customer's age, regressed on the full sample of prospective customers aged 18-79 (within support of 99% of the data). A 95% confidence interval band is shown in light grey. For our primary analyses, we denote customers aged 28 and younger (the median age for a credit card applicant) as less experienced, and customers 28 and older to be more experienced with credit cards. This cutoff is shown with a red vertical line, and it demonstrates that the probability a 28-year old prospective customer had a credit card at the time of this study was 58.5%.

Supplemental Summary Statistics Tables

In Table 1.1 of the paper, we present summary statistics for the control variables used in all regressions. Below, we present complementary summary statistics for the dependent measures. **Table A1.1** presents summary statistics for the full dataset. **Table A1.2** presents summary statistics during the non-promotion period. **Table A1.3** presents summary statistics during the promotion period. We have also included a correlation table in **Table A1.4**. All statistics represent each variable, presented at the equivalent level of analysis as the primary regressions presented and described in the paper.

The acquisition analysis sample (n=393,036) includes all prospective customers who arrived on the CBA credit card marketing website during the period of our study, minus noted exclusions (e.g., people who were assigned to both experimental conditions). The card application analysis sample (n=18,157) includes all prospective customers who applied for a credit card, so we could examine how the treatment affected the types of cards for which people applied. The engagement analyses sample (n=15,947) includes all customers who used their credit cards. Within the engagement analyses, we do look at different levels of data aggregation, depending on what best suits the hypothesis being tested:

- Spend analysis – Data are analyzed at the customer/month level, either during months where the customer still held their account (n=121,679) or with months after closing filled with zeroes, to determine no spend (n=126,658).

- Late payment analysis – Data are analyzed at the customer/month level, just including months for which a payment was due (n=112,641) or all months for which a payment was due within the first six months of having the card (n=92,988)

- Retention analysis – Data are analyzed at the customer level, either after six months (n=15,942) or after nine months (n=6,314). For the latter analysis, we only included customers in the analysis who opened their credit cards in time to have used them for nine months during the period of our study.

	All			**Control**			**Treatment**			**Diff.**
	N	Mean	SD	N	Mean	SD	N	Mean	SD	P-value
Dependent Variables										
App start	393,036	16.39%	0.37	195,438	16.66%	0.37	197,598	16.13%	0.37	0.00
Soft submission	393,036	13.31%	0.34	195,438	13.50%	0.34	197,598	13.12%	0.34	0.00
Submission	393,036	6.18%	0.24	195,438	6.20%	0.24	197,598	6.16%	0.24	0.59
Opened	393,036	4.62%	0.21	195,438	4.66%	0.21	197,598	4.58%	0.21	0.21
Activated	393,036	4.38%	0.20	195,438	4.43%	0.21	197,598	4.34%	0.20	0.15
Used	393,036	4.06%	0.20	195,438	4.09%	0.20	197,598	4.02%	0.20	0.24
Got Awards card	18,157	4.89%	0.22	9,111	4.80%	0.21	9,046	4.97%	0.22	0.58
Got Diamond card	18,157	2.21%	0.15	9,111	2.15%	0.15	9,046	2.28%	0.15	0.56
Got Gold card	18,157	2.76%	0.16	9,111	2.74%	0.16	9,046	2.77%	0.16	0.90
Got Low Fee card	18,157	14.67%	0.35	9,111	14.73%	0.35	9,046	14.61%	0.35	0.83
Got Low Fee Gold card	18,157	29.13%	0.17	9,111	27.55%	0.16	9,046	30.73%	0.17	0.20
Got Low Rate card	18,157	50.78%	0.50	9,111	51.84%	0.50	9,046	49.71%	0.50	0.00
Got Low Rate Gold card	18,157	12.46%	0.33	9,111	12.00%	0.32	9,046	12.92%	0.34	0.06
Got Platinum card	18,157	5.72%	0.23	9,111	5.42%	0.23	9,046	6.01%	0.24	0.09
Got Student card	18,157	3.60%	0.19	9,111	3.57%	0.19	9,046	3.64%	0.19	0.80
Monthly spend (Missing=.)	121,714	$ 187.52	17.17	60,918	$ 186.16	17.25	60,796	$ 188.89	17.08	0.37
Monthly spend (Missing=0)	126,696	$ 152.64	19.42	63,466	$ 150.92	19.55	63,230	$ 154.38	19.29	0.17
Open (6 months)	15,947	93.61%	0.24	8,003	93.53%	0.25	7,944	93.69%	0.24	0.67
Open (9 months)	6,314	90.51%	0.29	3,137	89.99%	0.30	3,177	91.03%	0.29	0.16
Late pay (6 months)	93,016	7.03%	0.26	46,674	7.15%	0.26	46,342	6.90%	0.25	0.14
Late pay (9 months)	112,674	7.51%	0.26	56,455	7.61%	0.27	56,219	7.40%	0.26	0.18

Table A1.1. Summary statistics of the dependent variables for customers in the control and treatment conditions (full sample).

	All			Control			Treatment			Diff.
	N	Mean	SD	N	Mean	SD	N	Mean	SD	P-value
Dependent Variables										
App start	145,412	14.46%	0.3516526	72,521	14.64%	0.35	72,891	14.27%	0.35	0.05
Soft submission	145,412	11.56%	0.32	72,521	11.64%	0.32	72,891	11.49%	0.32	0.38
Submission	145,412	4.64%	0.21	72,521	4.62%	0.21	72,891	4.66%	0.21	0.72
Opened	145,412	3.50%	0.18	72,521	3.46%	0.18	72,891	3.54%	0.18	0.44
Activated	145,412	3.27%	0.18	72,521	3.25%	0.18	72,891	3.29%	0.18	0.69
Used	145,412	2.98%	0.17	72,521	2.97%	0.17	72,891	2.99%	0.17	0.78
Got Awards card	5,089	6.07%	0.24	2,511	6.13%	0.24	2,578	6.01%	0.24	0.86
Got Diamond card	5,089	2.95%	0.17	2,511	3.19%	0.18	2,578	2.72%	0.16	0.32
Got Gold card	5,089	3.48%	0.18	2,511	3.82%	0.19	2,578	3.14%	0.17	0.18
Got Low Fee card	5,089	18.94%	0.39	2,511	18.48%	0.39	2,578	19.39%	0.40	0.40
Got Low Fee Gold card	5,089	4.11%	0.20	2,511	4.10%	0.20	2,578	4.11%	0.20	0.99
Got Low Rate card	5,089	41.48%	0.49	2,511	43.13%	0.50	2,578	39.88%	0.49	0.02
Got Low Rate Gold card	5,089	11.57%	0.32	2,511	10.43%	0.31	2,578	12.68%	0.33	0.01
Got Platinum card	5,089	6.56%	0.25	2,511	5.77%	0.23	2,578	7.33%	0.26	0.03
Got Student card	5,089	4.83%	0.21	2,511	4.94%	0.22	2,578	4.73%	0.21	0.73
Monthly spend (Missing=.)	40,701	170.66	18.30	20,174	165.49	18.57	20,527	175.90	18.02	0.03
Monthly spend (Missing=0)	41,802	149.05	19.38	20,770	142.92	20.15	21,032	155.37	19.38	0.00
Open (6 months)	4,331	97.37%	0.16	2,151	97.02%	0.17	2,180	97.71%	0.15	0.16
Open (9 months)	4,194	93.09%	0.25	2,086	92.28%	0.27	2,108	93.88%	0.24	0.04
Late pay (6 months)	25,902	8.12%	0.27	12,871	8.51%	0.28	13,031	7.73%	0.27	0.02
Late pay (9 months)	38,080	8.93%	0.29	18,897	9.20%	0.29	19,183	8.66%	0.28	0.07

Table A1.2. Summary statistics of the dependent variables for customers in the control and treatment conditions (non-promotion period).

	All			Control			Treatment			Diff.
	N	Mean	SD	N	Mean	SD	N	Mean	SD	P-value
Dependent Variables										
App start	247,624	17.53%	0.38	122,917	17.86%	0.38	124,707	17.21%	0.38	0.00
Soft submission	247,624	14.33%	0.35	122,917	14.59%	0.35	124,707	14.07%	0.35	0.00
Submission	247,624	7.09%	0.26	122,917	7.14%	0.26	124,707	7.04%	0.26	0.34
Opened	247,624	5.28%	0.22	122,917	5.37%	0.23	124,707	5.19%	0.22	0.04
Activated	247,624	5.03%	0.22	122,917	5.12%	0.22	124,707	4.95%	0.22	0.04
Used	247,624	4.69%	0.21	122,917	4.76%	0.21	124,707	4.62%	0.21	0.10
Got Awards card	13,068	4.42%	0.21	6,600	4.29%	0.20	6,468	4.56%	0.21	0.45
Got Diamond card	13,068	1.93%	0.14	6,600	1.76%	0.13	6,468	2.10%	0.14	0.15
Got Gold card	13,068	2.48%	0.16	6,600	2.33%	0.15	6,468	2.63%	0.16	0.28
Got Low Fee card	13,068	13.01%	0.34	6,600	13.30%	0.34	6,468	12.71%	0.33	0.31
Got Low Fee Gold card	13,068	2.45%	0.15	6,600	2.24%	0.15	6,468	2.66%	0.16	0.12
Got Low Rate card	13,068	54.40%	0.50	6,600	55.15%	0.50	6,468	53.63%	0.50	0.08
Got Low Rate Gold card	13,068	12.80%	0.33	6,600	12.59%	0.33	6,468	13.02%	0.34	0.47
Got Platinum card	13,068	5.39%	0.23	6,600	5.29%	0.22	6,468	5.49%	0.23	0.61
Got Student card	13,068	3.12%	0.17	6,600	3.05%	0.17	6,468	3.20%	0.18	0.61
Monthly spend (Missing=.)	81,013	$ 196.61	16.60	40,744	$ 197.34	16.59	40,269	$ 195.87	16.60	0.71
Monthly spend (Missing=0)	84,894	$ 154.44	19.25	42,696	$ 154.98	19.25	42,198	$ 153.89	19.25	0.73
Open (6 months)	11,616	92.21%	0.27	5,852	92.24%	0.27	5,764	92.18%	0.27	0.89
Open (9 months)	2,120	85.42%	0.35	1,051	85.44%	0.35	1,069	85.41%	0.35	0.98
Late pay (6 months)	67,114	6.61%	0.25	33,803	6.63%	0.25	33,311	6.58%	0.25	0.79
Late pay (9 months)	74,594	6.78%	0.25	37,558	6.82%	0.25	37,036	6.75%	0.25	0.73

Table A1.3. Summary statistics of the dependent variables for customers in the control and treatment conditions (promotion period).

	Customer age	Customer tenure	Male indicator	Transaction product	Savings product	Home loan product	Home insurance policy	Personal loan product	Retirement product	Motor insurance policy	Term deposit product	Got Awards	Got Diamond
Customer age	1												
Customer tenure	0.55***	1											
Male indicator	0.0030*	0.040***	1										
Transaction product	-0.29***	-0.12***	0.0082***	1									
Savings product	-0.092***	-0.0034**	0.031***	0.16***	1								
Home loan product	0.14***	0.15***	-0.0052***	0.13***	-0.13***	1							
Home insurance policy	0.16***	0.15***	-0.0012	0.057***	-0.029***	0.37***	1						
Personal loan product	-0.097***	0.0054***	0.0015	0.098***	0.054***	-0.045***	-0.0024	1					
Retirement product	-0.14***	-0.16***	-0.037***	0.062***	0.064***	-0.071***	-0.036***	0.016***	1				
Motor insurance policy	0.069***	0.050***	-0.0014	0.038***	0.0040**	0.13***	0.29***	0.061***	-0.0029*	1			
Term deposit product	0.16***	0.12***	0.027***	-0.029***	0.078***	-0.019***	0.031***	-0.052***	-0.022***	0.0090***	1		
Awards	-0.027***	-0.017***	0.00049	0.013***	0.014***	-0.013***	-0.0079***	0.012***	0.011***	-0.0013	-0.0029*	1	
Diamond	-0.017***	-0.014***	-0.012***	0.0085***	0.00047	0.014***	0.0039**	0.0028*	0.0045***	0.00071	-0.0023	-0.0015	1
Gold	-0.021***	-0.012***	-0.0045***	0.0076***	0.012***	-0.0076***	-0.0055***	0.0054***	0.0031*	0.0021	-0.0017	-0.0017	-0.0011
Low Fee	-0.047***	-0.040***	0.0021	0.021***	0.017***	-0.025***	-0.018***	0.0080***	0.014***	-0.0059***	-0.0049***	-0.0039**	-0.0026*
Low Fee Gold	-0.021***	-0.024***	-0.0077***	0.0088***	0.0075***	-0.011***	-0.0076***	0.0039**	0.0071***	-0.00084	-0.0045***	-0.0017	-0.0012
Low Rate	-0.087***	-0.065***	0.0047***	0.035***	0.027***	-0.046***	-0.031***	0.061***	0.019***	-0.010***	-0.017***	-0.0074***	-0.0050***
Low Rate Gold	-0.034***	-0.034***	-0.0062***	0.020***	0.012***	-0.019***	-0.012***	0.033***	0.012***	-0.0025	-0.0099***	-0.0036**	-0.0024
Platinum	-0.023***	-0.020***	-0.013***	0.016***	0.0050***	-0.0037**	-0.0047***	0.033***	0.011***	0.002	-0.0086***	-0.0024	-0.0016
Student	-0.046***	-0.033***	0.0029*	0.013***	0.017***	-0.022***	-0.014***	-0.0074***	0.011***	-0.0065***	-0.0050***	-0.0019	-0.0013
Spend 6 month	0.062***	0.045***	-0.027***	0.022**	-0.012	0.13***	0.037***	-0.033***	-0.0058	0.019**	0.002	0.081***	0.23***
Spend 9 month	0.058***	0.040***	-0.034**	0.01	0.00039	0.13***	0.023	-0.050***	-0.0025	0.0066	0.0069	0.059***	0.23***
Retention 6 month	-0.018**	-0.032***	0.0061	0.11***	0.00011	-0.064***	-0.013*	0.057***	0.029***	0.0019	-0.037***	0.037***	0.024***
Retention 9 month	-0.0096	-0.032***	0.0065	0.095***	0.00024	-0.063***	-0.015*	0.052***	0.026***	0.0052	-0.038***	0.031***	0.021***
Late pay 6 month	-0.046***	-0.039***	-0.025***	0.015*	-0.054***	-0.050***	-0.020**	0.075***	0.013	0.0016	-0.030***	-0.027***	0.017**
Late pay 9 month	-0.062***	-0.045***	-0.0064	0.014	-0.024	-0.038**	-0.022	0.064***	0.018	-0.012	-0.014	-0.041**	0.0066

	Got Gold	Got Low Fee	Got Low Fee Gold	Got Low Rate	Got Low Rate Gold	Got Platinum	Got Student	Spend 6 month	Spend 9 month	Retention 6 month	Retention 9 month	Late pay 6 month	Late pay 9 month
Customer age													
Customer tenure													
Male indicator													
Transaction product													
Savings product													
Home loan product													
Home insurance policy													
Personal loan product													
Retirement product													
Motor insurance policy													
Term deposit product													
Awards													
Diamond													
Gold	1												
Low Fee	-0.0030*	1											
Low Fee Gold	-0.0013	-0.0030*	1										
Low Rate	-0.0055***	-0.013***	-0.0057***	1									
Low Rate Gold	-0.0027*	-0.0063***	-0.0028*	-0.012***	1								
Platinum	-0.0018	-0.0043***	-0.0019	-0.0080***	0.39**	1							
Student	-0.0015	-0.0034**	-0.0015	-0.0063***	0.31*	-0.0021	1						
Spend 6 month	0.100***	-0.027***	0.021**	-0.15***	0.79	0.077***	-0.040***	1					
Spend 9 month	0.063***	-0.026*	0.022	-0.12***	0.66	0.060***	-0.037***	0.99***	1				
Retention 6 month	0.023***	0.063***	0.026***	-0.15***	0.46***	0.042***	0.038***	0.091***	0.067***	1			
Retention 9 month	0.028***	0.056***	0.024***	-0.13***	0.36***	0.035***	0.029***	0.10***	0.085***	0.79***	1		
Late pay 6 month	-0.021**	-0.0055	-0.000054	0.0063	0.8	0.019**	-0.0089	-0.057***	-0.071***	0.074***	0.089***	1	
Late pay 9 month	0.00041	-0.041**	-0.0059	0.037**	0.85	0.015	0.0046	-0.042**	-0.051**	0.053***	0.077***	0.30***	1

Table A1.4. Correlation Matrix for Independent and Dependent Variables.

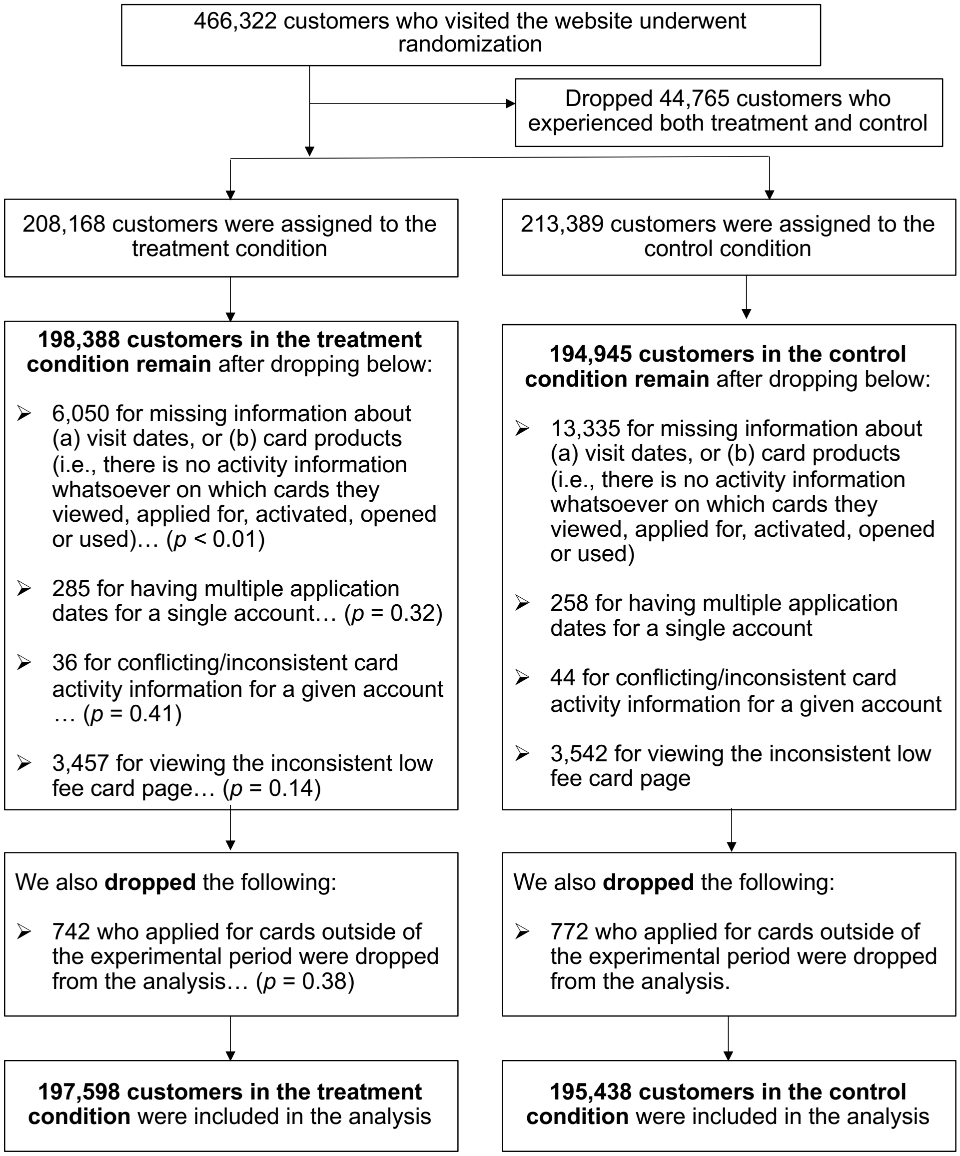

	All			Control			Treatment			Diff
	Missing	Total	% Missing	Missing	Total	% Missing	Missing	Total	% Missing	P-value
Customer age	94	393,036	0.02%	52	195,438	0.03	42	197,598	0.02	0.28
Customer tenure (months)	679	393,036	0.17%	346	195,438	0.18	333	197,598	0.17	0.52
Male indicator	173	393,036	0.04%	88	195,438	0.05	85	197,598	0.04	0.76
Transaction product	162	393,036	0.04%	83	195,438	0.04	79	197,598	0.04	0.70
Savings product	162	393,036	0.04%	83	195,438	0.04	79	197,598	0.04	0.70
Home loan product	162	393,036	0.04%	83	195,438	0.04	79	197,598	0.04	0.70
Home insurance policy	162	393,036	0.04%	83	195,438	0.04	79	197,598	0.04	0.70
Personal loan product	162	393,036	0.04%	83	195,438	0.04	79	197,598	0.04	0.70
Retirement product	162	393,036	0.04%	83	195,438	0.04	79	197,598	0.04	0.70
Motor insurance policy	162	393,036	0.04%	83	195,438	0.04	79	197,598	0.04	0.70
Term deposit product	162	393,036	0.04%	83	195,438	0.04	79	197,598	0.04	0.70

Figure A1.2. Randomization and Dropped Observations. P-values in parentheses indicate the statistical significance of the proportional difference in the number of dropped observations. For example, p=0.32 indicates that 258 dropped observations among 213,389 customers in the control condition is not significant from 285 dropped observations among 208,168 treated customers. Moreover, table reveals the number of control variables missing in the control and treatment conditions.

Supplemental Acquisition Analyses

In the primary analyses presented in the paper, we present the controlled analyses for all customers who visited the credit card website during the experimental period, as described in Section 3.2.1. Here, we present two additional analyses. First, **Figures A1.1** and **A1.2** present split-sample analyses for customers with low experience levels (ages 28 and under) and high experience levels (over 28 years of age). Consistent with the primary results presented, we find no differences in net acquisition rates among customers in the treatment and control condition. Second, **Figure A1.3** replicates the primary results without control variables. This is important since controlled analyses will exclude some customers who did not start the application process, owing to missing control variables. Again, we note that the results are substantively similar if control variables are excluded from the analysis.

Finally, **Tables A1.5-A1.8** present the marginal effects and significance values for the acquisition analyses presented in figures. Rather than presenting the results from the interaction model shown as Equation 1 in the paper, we formulate the analyses in **Tables A1.5-A1.8** from conducting split sample analyses during the promotion and non-promotion periods. The results obtained are equivalent to the fully-interacted model, but they afford a clearer picture of the relative number of observations collected during the non-promotion and promotion periods.

	No Promo					Promo				
Stage	n	Treatment/ no promotion	Control / no promotion	sig	p-val	n	Treatment/ promotion	Control/ promotion	sig	p-val
Application Started	145,030	14.55%	14.40%		0.491	245,244	17.33%	17.78%	***	0.001
Soft submission	145,030	11.67%	11.50%		0.375	245,244	14.17%	14.53%	***	0.007
Submission	142,319	4.73%	4.74%		0.979	245,244	7.08%	7.12%		0.711
Opened	145,030	3.54%	3.47%		0.436	245,244	5.22%	5.36%		0.122
Activated	145,030	3.29%	3.27%		0.876	245,244	4.98%	5.11%		0.122
Used	144,843	2.99%	2.99%		0.963	245,244	4.65%	4.75%		0.221

Table A1.5. Marginal Effects Results Table Corresponding with the Acquisition Funnel presented in Figure 4 in the manuscript. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively.

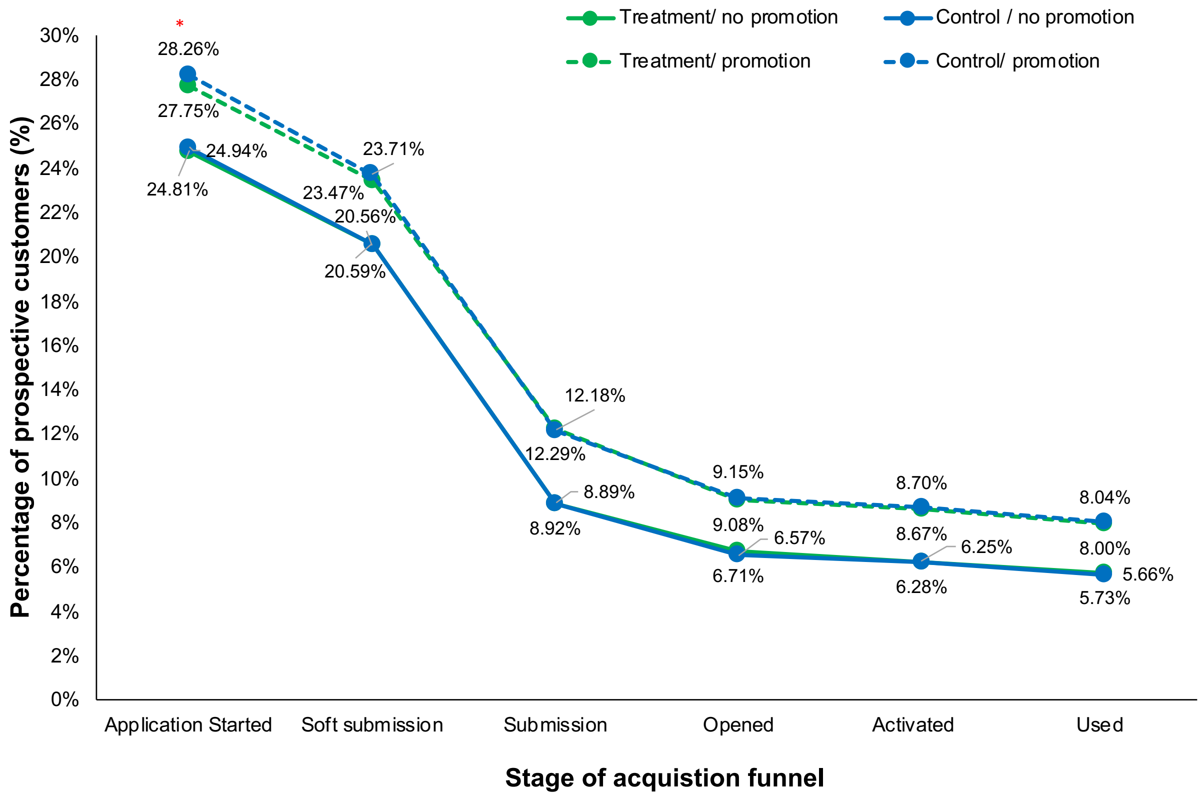

Figure A1.1: Acquisition funnel conversion percentages during the promotion and non-promotion periods for less experienced customers (28 years of age and younger). Marginal effects are from logistic regression models, estimated with robust standard errors clustered at the first website visit date level. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively. The results indicate that, consistent with Figure 1.5, conversion rates are significantly higher during the promotion period. Moreover, although fewer customers start the application process in the treatment condition during the promotion period, the treatment ultimately had no effect on rates of customer acquisition, neither during the promotion nor the non-promotion periods.

	No Promo					Promo				
Stage	**n**	**Treatment/ no promotion**	**Control / no promotion**	**sig**	**p-val**	**n**	**Treatment/ promotion**	**Control/ promotion**	**sig**	**p-val**
Application Started	39,661	24.81%	24.94%		0.760	70,831	27.75%	28.26%	*	0.078
Soft submission	39,661	20.59%	20.56%		0.956	70,831	23.47%	23.71%		0.371
Submission	38,661	8.92%	8.89%		0.914	70,831	12.29%	12.18%		0.609
Opened	39,661	6.71%	6.57%		0.540	70,831	9.08%	9.15%		0.732
Activated	39,538	6.28%	6.25%		0.878	70,831	8.67%	8.70%		0.894
Used	39,538	5.73%	5.66%		0.728	70,831	8.00%	8.04%		0.860

Table A1.6. Marginal Effects Results Table Corresponding with the Acquisition Funnel presented in Figure A1.1 in the Appendix. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively.

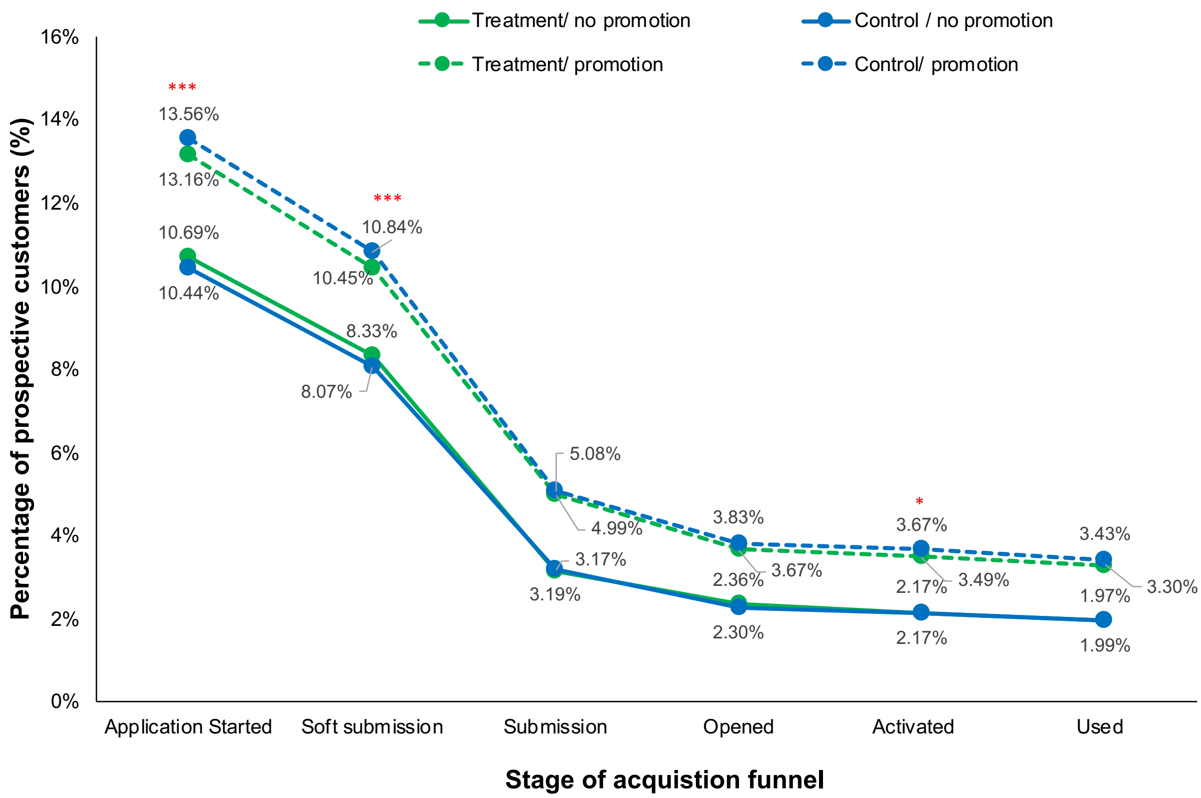

Figure A1.2: Acquisition funnel conversion percentages during the promotion and non-promotion periods for more experienced customers (over 28 years of age). Marginal effects are from logistic regression models, estimated with robust standard errors clustered at the first website visit date level. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively. The results indicate that, consistent with Figure 1.5, conversion rates are significantly higher during the promotion period. Moreover, although fewer customers start the application process or soft submit their application in the treatment condition during the promotion period, the treatment ultimately had no effect on rates of customer acquisition, neither during the promotion nor the non-promotion periods.

	No Promo					Promo				
Stage	**n**	**Treatment/ no promotion**	**Control / no promotion**	**sig**	**p-val**	**n**	**Treatment/ promotion**	**Control/ promotion**	**sig**	**p-val**
Application Started	105,369	10.69%	10.44%		0.267	176,276	13.16%	13.56%	***	0.009
Soft submission	105,369	8.33%	8.07%		0.167	176,276	10.45%	10.84%	***	0.007
Submission	103,658	3.17%	3.19%		0.900	176,276	4.99%	5.08%		0.405
Opened	105,369	2.36%	2.30%		0.542	176,276	3.67%	3.83%		0.104
Activated	105,040	2.17%	2.17%		0.957	176,276	3.49%	3.67%	*	0.064
Used	104,911	1.97%	1.99%		0.787	176,276	3.30%	3.43%		0.160

Table A1.7. Marginal Effects Results Table Corresponding with the Acquisition Funnel presented in Figure A1.2 in the Appendix. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively.

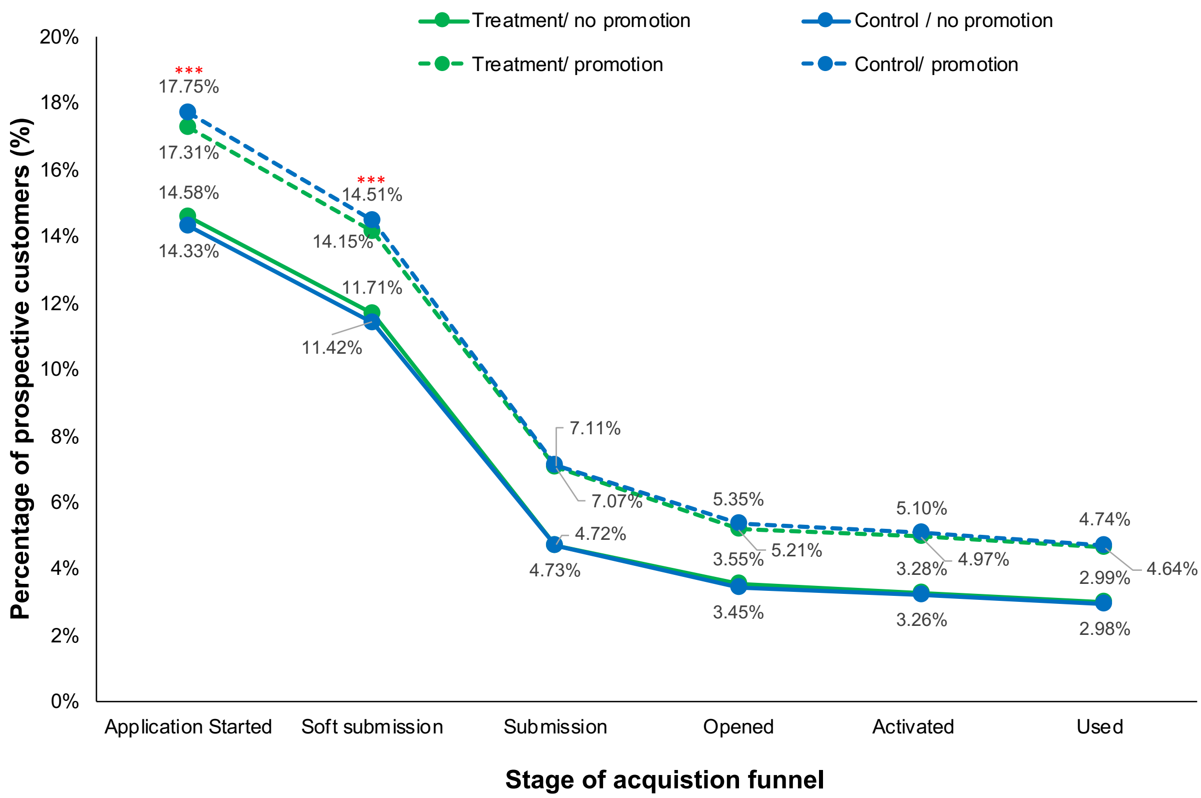

Figure A1.3: Acquisition funnel conversion percentages during the promotion (n=245,756) and non-promotion (n=145,412) periods excluding control variables. Date fixed effects are retained. This analysis enables the inclusion of observations for which control variables were missing. Beyond the acquisition phase, prospective customers were obliged to provide information about themselves, which entered into the controlled analyses. Marginal effects are from logistic regression models, estimated with robust standard errors clustered at the first website visit date level. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively. The results indicate that conversion rates are significantly higher during the promotion period. Moreover, although fewer customers who are exposed to the treatment during the promotion period start or soft submit an application, the treatment ultimately had no effect on rates of customer acquisition, neither during the promotion nor the non-promotion periods.

	No Promo					Promo				
Stage	**n**	**Treatment/ no promotion**	**Control / no promotion**	**sig**	**p-val**	**n**	**Treatment/ promotion**	**Control/ promotion**	**sig**	**p-val**
Application Started	145,412	14.58%	14.33%		0.297	245,756	17.31%	17.75%	***	0.001
Soft submission	145,412	11.71%	11.42%		0.168	245,756	14.15%	14.51%	***	0.006
Submission	142,694	4.72%	4.73%		0.959	245,756	7.07%	7.11%		0.702
Opened	145,412	3.55%	3.45%		0.343	245,756	5.21%	5.35%		0.118
Activated	145,412	3.28%	3.26%		0.833	245,756	4.97%	5.10%		0.116
Used	145,224	2.99%	2.98%		0.928	245,756	4.64%	4.74%		0.215

Table A1.8. Marginal Effects Results Table Corresponding with the Acquisition Funnel presented in Figure A1.3 in the Appendix. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively.

Supplemental Card Choice Analyses

In the primary analysis presented in the paper, **Table 1.2** shows the marginal effects that emerge from the results of the logistic regression models specified in **Equation (2)** in the paper and presented in regression table form in **Table A1.9**. The results demonstrate that, consistent with H2, customers make different choices in the presence of tradeoff transparency.

	(1) Awards	(2) Diamond	(3) Gold	(4) Low Fee	(5) Low Fee Gold	(6) Low Rate	(7) Low Rate Gold	(8) Platinum	(9) Student
Treatment	-0.014	-0.133	-0.193	0.064	0.005	-0.144**	0.212**	0.266**	-0.011
	(0.125)	(0.157)	(0.149)	(0.064)	(0.166)	(0.067)	(0.086)	(0.108)	(0.136)
Promotion	-0.150	-0.826***	-0.437**	-0.517***	-0.593***	0.498***	0.170	-0.084	-0.325**
	(0.130)	(0.273)	(0.179)	(0.087)	(0.175)	(0.062)	(0.107)	(0.136)	(0.154)
Treatment x promotion	0.075	0.286	0.310*	-0.126	0.170	0.088	-0.162	-0.221	0.061
	(0.150)	(0.210)	(0.188)	(0.084)	(0.206)	(0.077)	(0.101)	(0.137)	(0.173)
Customer age	0.024	0.060	0.000	-0.035***	-0.002	0.035***	0.039***	0.116***	-0.400***
	(0.023)	(0.037)	(0.027)	(0.011)	(0.022)	(0.009)	(0.012)	(0.021)	(0.027)
Customer age^2	-0.000	-0.001**	-0.000	0.000***	0.000	-0.000***	-0.000	-0.001***	0.003***
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Customer tenure	0.004***	-0.005***	0.002	-0.001*	-0.004***	0.002***	-0.001	-0.000	-0.008***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
Customer tenure^2	-0.000**	0.000***	-0.000	0.000	0.000	-0.000***	0.000	0.000	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.027	-0.724***	-0.280***	0.097**	-0.373***	0.216***	-0.135***	-0.498***	0.202**
	(0.069)	(0.123)	(0.096)	(0.038)	(0.092)	(0.033)	(0.047)	(0.065)	(0.080)
Retirement product	0.263**	0.061	-0.201	-0.010	-0.047	-0.090	0.080	0.236**	-0.124
	(0.115)	(0.180)	(0.179)	(0.073)	(0.172)	(0.058)	(0.085)	(0.107)	(0.144)
Home loan product	-0.099	1.794***	0.233	-0.106	-0.101	-0.299***	-0.075	0.510***	-2.804***
	(0.118)	(0.128)	(0.152)	(0.081)	(0.177)	(0.052)	(0.088)	(0.110)	(1.005)
Personal loan product	-0.257***	-0.456***	-0.470***	-0.576***	-0.419***	0.295***	0.235***	0.518***	-1.487***
	(0.097)	(0.147)	(0.144)	(0.068)	(0.128)	(0.045)	(0.052)	(0.081)	(0.171)
Savings product	0.296***	-0.030	0.445***	0.034	0.100	-0.112***	0.002	-0.073	0.261**
	(0.083)	(0.110)	(0.132)	(0.051)	(0.097)	(0.038)	(0.059)	(0.075)	(0.130)
Term deposit product	0.309	0.507*	0.433	0.457***	-0.148	-0.159	-0.412**	-0.942***	0.197
	(0.224)	(0.302)	(0.285)	(0.148)	(0.347)	(0.099)	(0.172)	(0.365)	(0.343)
Transaction product	0.169	-0.162	-0.447*	0.167	0.010	-0.438***	0.405***	0.928***	1.017*
	(0.240)	(0.364)	(0.245)	(0.127)	(0.267)	(0.091)	(0.140)	(0.305)	(0.597)
Home insurance policy	0.114	0.342*	-0.035	-0.144	-0.072	-0.040	0.077	0.056	-1.381**
	(0.147)	(0.201)	(0.197)	(0.089)	(0.250)	(0.067)	(0.108)	(0.138)	(0.700)
Motor insurance policy	0.122	-0.175	0.614***	-0.042	0.258	-0.152*	-0.037	0.174	-0.063
	(0.208)	(0.283)	(0.206)	(0.135)	(0.258)	(0.082)	(0.127)	(0.171)	(0.413)
Constant	-3.724***	-3.251***	-2.830***	-0.810***	-2.639***	-0.746***	-3.327***	-5.929***	3.941***
	(0.498)	(0.872)	(0.583)	(0.283)	(0.495)	(0.205)	(0.277)	(0.578)	(0.662)
Observations	18,152	18,152	18,152	18,152	18,152	18,152	18,152	18,152	18,152
Sample	All	All	All	All	All	All	All	All	All
R-squared	0.200	0.200	0.200	0.200	0.200	0.200	0.200	0.200	0.200

Table A1.9: Effects of the treatment on the credit card choices of customers. Columns (1)-(9) show the results of the logistic regression models with robust standard errors clustered by customers' first website visit dates. All models additionally include indicator variables for the week during which the customer first visited the credit card marketing website, which are withheld from the table for parsimony. Indicator variables for *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively. Low rate and low rate gold cards were part of the publicly announced promotion campaign beginning on November 1, 2017. The results indicate that providing transparency into the tradeoffs of the service offering influenced credit card choices, which is consistent with H2.

An alternative question is whether tradeoff transparency affected customers' choices for the better. In the paper, we present results demonstrating that customers use the cards more, are less likely to make late payments, and are more likely to keep the cards longer, if they are provided with tradeoff transparency when they are choosing credit cards. These results serve as our primary evidence that tradeoff transparency affects customers' choices for the better, since these are behavioral indicators that suggest that customers find the credit cards they choose in the presence of tradeoff transparency to be more valuable, and less burdensome to their financial health. An alternative way of assessing the quality of customer choices is to examine whether, based on observable customer-level characteristics, the cards customers select in the presence of tradeoff transparency appear to be in better alignment with their presumed financial needs than the cards customers select in the absence of tradeoff transparency.

In this supplemental analysis, we considered two observable demographic variables – customer age and total count of financial products. We presume that:

1. The student card might be disproportionately more aligned with the needs and preferences of younger customers, who are more likely to be students and have basic financial needs.

2. The low rate card might be disproportionately more aligned with the needs and preferences of customers with a less-developed financial position, who have fewer banking products, and who in turn may wish to use their credit card for borrowing, carrying a balance from month to month.

To test these hypotheses, we consider the interaction of the treatment variable with continuous variables measuring the customer's age, and their total count of financial products at the time they first visited the bank's credit card marketing website. In particular, we use separate regression specifications for analyzing the moderating effects of tradeoff transparency on influencing the choices of customers who vary in age and the breadth of their financial relationship with the bank.

In the first specification presented below, β_3 captures the degree to which tradeoff transparency has a differential effect on older customers to choose the student card.

$$\Pr\left(STUDENT_i\right) = f\left(\begin{array}{c} \beta_0 + \beta_1 TREAT_i + \beta_2 AGE + \beta_3 TREAT_i \times AGE_i + \\ \beta_4 TENURE_i + \beta_5 TENURE_i^2 + \beta_6 GENDER_i + \\ \beta_7 HLOAN_i + \beta_8 PLOAN_i + \beta_9 TRANS_i + \beta_{10} SAV_i + \beta_{11} HINS_i + \\ \beta_{12} VINS_i + \beta_{13} RET_i + X_i + \epsilon_i \end{array}\right) \tag{2}$$

If tradeoff transparency increases the probability that younger customers will select the student card, as predicted above, then β_3 will be negative and significant. In the second specification, γ_3 captures the degree to which tradeoff transparency has a differential effect on customers with broader financial relationships to select the low rate card.

$$\Pr\left(LOWRATE_i\right) = f\left(\begin{array}{c} \gamma_0 + \gamma_1 TREAT_i + \gamma_2 PRODUCTS_i + \gamma_3 TREAT_i \times PRODUCTS_i + \\ \gamma_4 AGE_i + \gamma_5 AGE_i^2 + \gamma_6 TENURE_i + \gamma_7 TENURE_i^2 + \gamma_8 GENDER_i + \\ \gamma_9 HLOAN_i + \gamma_{10} PLOAN_i + \gamma_{11} TRANS_i + \gamma_{12} SAV_i + \gamma_{13} HINS_i + \\ \gamma_{14} VINS_i + \gamma_{15} RET_i + X_i + \epsilon_i \end{array}\right) \tag{3}$$

If tradeoff transparency increases the probability that customers with narrower financial relationships with the bank will select the low rate card, then we would predict γ_3 will be negative and significant.

	(1)	(2)	(3)	(4)	(5)	(6)
	Pr(Student)	Pr(Student)	Pr(Student)	Pr(Low rate)	Pr(Low rate)	Pr(Low rate)
Treatment	-0.163	0.295	-0.363	0.067	0.257	-0.014
	(0.603)	(1.003)	(0.779)	(0.085)	(0.180)	(0.094)
Customer age	-0.225***	-0.219***	-0.226***	0.037***	0.001	0.047***
	(0.018)	(0.030)	(0.023)	(0.009)	(0.017)	(0.010)
Treatment x customer age	0.008	-0.013	0.018			
	(0.026)	(0.044)	(0.033)			
Product count				-0.086	-0.434	-0.059
				(0.100)	(0.278)	(0.113)
Treatment x product count				-0.069*	-0.179***	-0.021
				(0.035)	(0.068)	(0.039)
Customer age^2				-0.000***	-0.000	-0.001***
				(0.000)	(0.000)	(0.000)
Tenure	-0.008***	-0.007***	-0.009***	0.002***	0.000	0.002***
	(0.002)	(0.002)	(0.002)	(0.000)	(0.001)	(0.001)
Tenure^2	0.000***	0.000**	0.000***	-0.000***	0.000	-0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Gender	0.203**	0.257*	0.169*	0.204***	0.259***	0.197***
	(0.079)	(0.136)	(0.099)	(0.033)	(0.069)	(0.038)
Home loan indicator	-2.806***	-1.611		-0.169	0.263	-0.239*
	(1.004)	(1.003)		(0.114)	(0.295)	(0.132)
Home insurance indicator	-1.371*		-0.815	0.073	0.503*	0.021
	(0.702)		(0.692)	(0.114)	(0.298)	(0.129)
Motor insurance indicator	-0.054	0.210	-0.269	-0.038	0.435	-0.093
	(0.410)	(0.614)	(0.574)	(0.138)	(0.357)	(0.156)
Personal loan indicator	-1.449***	-1.681***	-1.342***	0.374***	1.174***	0.201*
	(0.170)	(0.266)	(0.219)	(0.107)	(0.277)	(0.119)
Transaction product indicator	1.017*		0.681	-0.347***	0.604*	-0.532***
	(0.585)		(0.585)	(0.132)	(0.334)	(0.146)
Savings product indicator	0.259**	0.465**	0.141	0.012	0.322	-0.018
	(0.126)	(0.217)	(0.156)	(0.103)	(0.265)	(0.119)
Retirement product indicator	-0.102	-0.062	-0.164	0.011	0.347	0.010
	(0.143)	(0.244)	(0.174)	(0.120)	(0.312)	(0.140)
Term deposit indicator	0.163	0.534	0.043			
	(0.344)	(0.679)	(0.399)			
Constant	1.751***	2.701***	2.098***	-0.292	-0.611*	-0.177
	(0.663)	(0.732)	(0.722)	(0.187)	(0.325)	(0.227)
Sample	All	Non-promo	Promo	All	Non-promo	Promo
Observations	18,152	4,678	11,632	18,152	5,087	13,065

Table A1.10: Moderating effects of tradeoff transparency on customers who vary by age and count of financial products. Models estimated with logistic regression with robust standard errors clustered by customers' first website visit dates. All models additionally include indicator variables for the week during which the customer first visited the credit card marketing website, which are withheld from the table for parsimony. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively. Models are estimated on the full sample of observations, and on subsamples of customer observations for customers who applied outside and inside the promotion period. Due to the logistic regression models, observations for which results were perfectly predicted by particular covariates were dropped. The results provide evidence that customers with fewer financial products were more likely to select the Low Rate card in the presence of tradeoff transparency, especially when a promotion was not taking place. Although younger customers were more likely to select the Student card, their propensity to do so did not depend upon the presence of tradeoff transparency.

As shown in **Table A1.10**, we do not find that younger customers are disproportionately likely to select the student card in the presence of tradeoff transparency, but we do find that customers with fewer financial products are more likely to select the low rate card in the presence of tradeoff transparency, especially during the non-promotion period.

Our intuition is that the student card is much more clearly branded than the other two cards as being "for students," and hence, the addition of tradeoff transparency has no effect on its relative appeal to younger customers. It may also be the case that since we find that tradeoff transparency is more effective for older customers who have more prior experience with credit cards, younger customers didn't get as much benefit from the tradeoff messaging. Nevertheless, the fact that the predicted results hold for customers with fewer products is directionally consistent with the idea that tradeoff transparency may help customers make *better* choices. We note these results are also consistent with the results of our tests of H3 and H4.

Finally, consistent with prior research demonstrating how customer compatibility can lead to more satisfying and profitable customer relationships, it is our account that through its capacity to inform customer decisions, tradeoff transparency drives profitability by enhancing customer engagement. However, an alternative account is that tradeoff transparency may lead companies to attract more profitable customers in general. Since all of the primary analyses presented in this paper control for observable customer characteristics, the results are consistent with the idea that otherwise-similar customers will exhibit higher levels of long-term engagement with they are provided with tradeoff transparency as a part of the customer acquisition process. To test whether tradeoff transparency may also lead to the attraction of more or less profitable customers, we analyze all customers who opened credit card accounts during the experiment, analyzing how the observable characteristics (age, tenure, and product ownership) differ among customers who were motivated to open accounts in the treatment and control conditions.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Age	Tenure	Male	Home loan	Home insurance	Motor insurance	Personal loan	Transaction product	Savings product	Retirement product	Term deposit
Treatment	-0.132	-3.799	0.057	-0.038	0.045	-0.003	0.092	0.217	0.012	0.176*	0.305
	(0.254)	(3.039)	(0.057)	(0.076)	(0.136)	(0.159)	(0.059)	(0.189)	(0.060)	(0.103)	(0.263)
Promo	0.575**	-4.681*	-0.031	-0.015	-0.094	-0.182	-0.254***	-0.297**	-0.016	-0.040	0.418**
	(0.233)	(2.686)	(0.045)	(0.075)	(0.120)	(0.150)	(0.065)	(0.141)	(0.056)	(0.088)	(0.206)
Treatment x promo	-0.401	1.831	-0.033	0.106	0.083	0.083	-0.167**	0.029	-0.019	-0.239**	-0.138
	(0.330)	(3.699)	(0.065)	(0.097)	(0.164)	(0.187)	(0.076)	(0.212)	(0.077)	(0.121)	(0.291)
Constant	31.107***	141.064***	-0.234***	-2.112***	-2.822***	-3.364***	-1.221***	3.628***	1.192***	-2.426***	-4.382***
	(0.182)	(2.191)	(0.039)	(0.062)	(0.097)	(0.129)	(0.054)	(0.123)	(0.045)	(0.075)	(0.180)
Observations	18,157	18,152	18,157	18,157	18,157	18,157	18,157	18,157	18,157	18,157	18,157
Model	OLS	OLS	Logistic	Logistic	Logistic	Logistic	Logistic	Logistic	Logistic	Logistic	Logistic
R-squared (Pseudo R2)	0.001	0.000	0.000145	0.000174	0.000441	0.000568	0.00412	0.00316	2.55e-05	0.00114	0.00309

Table A1.11: Comparison of Characteristics of Customers Attracted in the Presence and Absence of Tradeoff Transparency. Columns (1-2) are estimated with OLS regression and Columns (3-11) are estimated with logistic regression, all with robust standard errors, clustered by date of first website visit, shown in parentheses. models additionally include indicator variables for the week during which the customer first visited the credit card marketing website, which are withheld from the table for parsimony. *, **, and ***, signify significance at the 10%, 5%, and 1% levels respectively.

As demonstrated in **Table A1.11**, we observe that the bank attracted customers that were insignificantly different in the presence and absence of tradeoff transparency, with the exception that customers attracted in the presence of tradeoff transparency were marginally more likely to have a retirement product – the baseline probability increasing from 8.12% to 9.54% (β=0.176, p<0.10). However, we do observe that in the presence of a promotion, customers attracted tended to be older (β=0.575, p<0.05) and marginally newer to the bank (β=-4.681, p<0.10), and were less likely to have a personal loan account (β=-0.254, p<0.01 a transaction product (β =-0.297, p<0.01), but were more likely to have a term deposit product (β =0.418, p<0.05).

Customer Engagement Robustness Analysis: Considering Alternative Experience Delineations

In the primary analysis presented in the manuscript, a median split on age is used to distinguish customers with more and less prior credit card experience. **Tables A1.12-1.13** demonstrate that the results are substantively similar if the 25th percentile of age (24 years old) is instead used as the delineator.

	(1) Ln(Spend)	(2) Ln(Spend)	(3) Ln(Spend)	(4) Ln(Spend)	(5) Ln(Spend)	(6) Ln(Spend)
Treatment	0.080	-0.008	0.128**	0.095*	0.007	0.149***
	(0.052)	(0.138)	(0.060)	(0.054)	(0.144)	(0.057)
Promotion	-0.104	0.209	-0.251**	-0.371***	0.107	-0.588***
	(0.103)	(0.185)	(0.118)	(0.114)	(0.172)	(0.129)
Treatment x promotion	-0.097	-0.034	-0.138*	-0.109*	-0.039	-0.160**
	(0.061)	(0.153)	(0.071)	(0.064)	(0.158)	(0.070)
Customer age	0.072***	0.556	0.050***	0.046***	0.646	0.042***
	(0.008)	(0.408)	(0.014)	(0.009)	(0.409)	(0.014)
Customer age²	-0.001***	-0.012	-0.001***	-0.001***	-0.014	-0.000***
	(0.000)	(0.010)	(0.000)	(0.000)	(0.010)	(0.000)
Customer tenure	0.002***	0.006***	0.000	0.000	0.005***	-0.001**
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Customer tenure²	-0.000	-0.000***	0.000	0.000	-0.000**	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.050	0.069	0.031	0.072**	0.081	0.055
	(0.034)	(0.058)	(0.037)	(0.034)	(0.059)	(0.039)
Retirement product	0.067	-0.010	0.123	0.105*	-0.003	0.182**
	(0.058)	(0.095)	(0.076)	(0.058)	(0.096)	(0.077)
Home loan product	0.509***	0.546**	0.516***	0.349***	0.209	0.370***
	(0.049)	(0.219)	(0.053)	(0.062)	(0.229)	(0.063)
Personal loan product	-0.324***	-0.365***	-0.308***	-0.209***	-0.369***	-0.151***
	(0.039)	(0.082)	(0.048)	(0.043)	(0.084)	(0.052)
Savings product	0.277***	0.402***	0.235***	0.247***	0.376***	0.208***
	(0.040)	(0.084)	(0.043)	(0.040)	(0.080)	(0.043)
Term deposit product	0.198*	0.558***	0.100	-0.023	0.503**	-0.150
	(0.120)	(0.196)	(0.136)	(0.133)	(0.218)	(0.151)
Transaction product	0.496***	0.105	0.527***	1.043***	0.990**	1.026***
	(0.105)	(0.445)	(0.111)	(0.101)	(0.455)	(0.108)
Home insurance policy	-0.038	0.158	-0.044	-0.044	0.332	-0.070
	(0.076)	(0.235)	(0.080)	(0.085)	(0.235)	(0.091)
Motor insurance policy	0.045	0.189	0.044	0.038	0.151	0.045
	(0.105)	(0.239)	(0.114)	(0.108)	(0.245)	(0.117)
Constant	3.348***	-2.057	4.010***	3.455***	-3.696	3.843***
	(0.212)	(4.371)	(0.304)	(0.245)	(4.311)	(0.331)
Observations	121,679	36,540	85,139	126,658	37,478	89,180
Customers	15,942	4,686	11,256	15,942	4,686	11,256
Data treatment for closed accounts	Missing	Missing	Missing	Zero	Zero	Zero
Sample	All	Age≤24	Age>24	All	Age≤24	Age>24
R-squared	0.0321	0.0463	0.0277	0.0271	0.0403	0.0292
Pred(Y): Non-Promotion: Control	$196.64	$132.62	$235.76	$192.67	$124.92	$235.17
Pred(Y): Non-Promotion: Treatment	$212.98	$131.62	$268.03	$211.79	$125.79	$272.87
Pred(Y): Promotion: Control	$177.20	$163.51	$183.48	$132.91	$139.02	$130.63
Pred(Y): Promotion: Treatment	$174.20	$156.93	$181.69	$130.99	$134.57	$129.23

Table A1.12: Effects of the treatment on monthly spend with alternative age cutoff (24 years). Columns (1) – (3) show the results of panel data regression models for all customers as well as younger and older customers, with spend values for months after cancellation set to missing. Columns (4) – (6) show the same specifications with spend values for months after cancellation set to zero. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the bottom section of the table. Consistent with H3, H5, and H7, the results indicate that customers exposed to the treatment spent more on their cards, that these effects were stronger for more experienced customers, and that these effects were attenuated by the promotion.

	(1) Pr(Retain6)	(2) Pr(Retain6)	(3) Pr(Retain6)	(4) Pr(Retain9)	(5) Pr(Retain9)	(6) Pr(Retain9)
Treatment	0.259	0.234	0.314	0.249**	0.247	0.274**
	(0.202)	(0.389)	(0.208)	(0.115)	(0.212)	(0.131)
Promotion	-1.736***	-1.321*	-1.807***	-1.344***	-0.277	-1.658***
	(0.256)	(0.758)	(0.248)	(0.231)	(0.490)	(0.237)
Treatment x promotion	-0.239	-0.115	-0.313	-0.223	-0.033	-0.306
	(0.220)	(0.426)	(0.229)	(0.165)	(0.293)	(0.186)
Customer age	-0.219***	0.634	-0.072*	-0.094**	-0.573	-0.065
	(0.038)	(0.975)	(0.042)	(0.041)	(1.078)	(0.053)
Customer age^2	0.003***	-0.018	0.001**	0.001***	0.015	0.001*
	(0.001)	(0.023)	(0.001)	(0.001)	(0.026)	(0.001)
Customer tenure	-0.009***	-0.011***	-0.009***	-0.005***	-0.007	-0.005***
	(0.001)	(0.004)	(0.001)	(0.002)	(0.005)	(0.002)
Customer tenure^2	0.000***	0.000*	0.000***	0.000***	0.000	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.046	0.042	0.049	0.078	0.288*	0.005
	(0.074)	(0.157)	(0.080)	(0.091)	(0.168)	(0.123)
Retirement product	0.245	0.131	0.284*	-0.107	-0.418	0.036
	(0.154)	(0.322)	(0.162)	(0.155)	(0.258)	(0.194)
Home loan product	-0.595***	-1.322***	-0.556***	-0.564***	-1.118**	-0.534***
	(0.105)	(0.334)	(0.111)	(0.167)	(0.464)	(0.175)
Personal loan product	0.767***	-0.005	0.968***	0.246**	-0.250	0.425***
	(0.110)	(0.205)	(0.142)	(0.106)	(0.236)	(0.145)
Savings product	-0.188**	-0.173	-0.163**	-0.132	0.160	-0.194*
	(0.080)	(0.231)	(0.078)	(0.096)	(0.203)	(0.103)
Term deposit product	-0.617***	-0.122	-0.658***	-0.662***	-0.383	-0.657**
	(0.167)	(0.621)	(0.181)	(0.235)	(0.637)	(0.265)
Transaction product	1.603***	2.696***	1.468***	1.402***	1.976***	1.344***
	(0.125)	(0.494)	(0.133)	(0.222)	(0.641)	(0.254)
Home insurance policy	0.084	1.709*	0.012	-0.068		-0.206
	(0.154)	(1.035)	(0.162)	(0.170)		(0.179)
Motor insurance policy	0.082	0.103	0.087	0.027	0.504	0.058
	(0.198)	(0.630)	(0.209)	(0.202)	(0.932)	(0.222)
Constant	7.509***	-2.371	4.791***	3.739***	7.499	3.145***
	(0.711)	(10.550)	(0.809)	(0.773)	(11.368)	(1.068)
Observations	15,942	4,278	11,256	6,314	1,924	4,357
Customers	All	Age≤24	Age>24	All	Age≤24	Age>24
Pseudo R2	0.0845	0.0666	0.0883	0.0551	0.0462	0.0760
Pred(Y): Non-Promotion: Control	98.07%	98.19%	97.90%	98.07%	98.19%	97.90%
Pred(Y): Non-Promotion: Treatment	98.50%	98.56%	98.45%	98.50%	98.56%	98.45%
Pred(Y): Promotion: Control	90.50%	93.87%	89.07%	90.50%	93.87%	89.07%
Pred(Y): Promotion: Treatment	90.67%	94.49%	89.08%	90.67%	94.49%	89.08%

Table A1.13: Effects of the treatment on customer retention rates with alternative age cutoff (24 years). Effects of the treatment on customer retention rates. Columns (1) – (3) and columns (4) – (6) show the results of cross-sectional data regression models of all, younger, and older customers for retention six and nine months after card activation, respectively. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the bottom section of the table. The results indicate that transparency can lead to increased customer retention, which is consistent with H4, that the effects are strongest for more experienced customers, which is consistent with H6, and that the effects are dampened during the promotion, which is consistent with H8.

Customer Engagement Robustness Analysis: Inconsistent Benefits Information

In the primary analysis presented in the manuscript, we excluded observations from customers who observed inconsistent benefits information during their deliberation process. **Tables A1.14-15** reveal that the results are substantively similar if those observations are included in the analysis.

	(1) Ln(Spend)	(2) Ln(Spend)	(3) Ln(Spend)	(4) Ln(Spend)	(5) Ln(Spend)	(6) Ln(Spend)
Treatment	0.093*	0.051	0.148**	0.108**	0.047	0.190***
	(0.051)	(0.088)	(0.061)	(0.053)	(0.094)	(0.054)
Promotion	-0.083	0.026	-0.194	-0.348***	-0.139	-0.556***
	(0.099)	(0.131)	(0.146)	(0.111)	(0.125)	(0.162)
Treatment x promotion	-0.107*	-0.076	-0.159**	-0.121*	-0.070	-0.198***
	(0.060)	(0.104)	(0.075)	(0.062)	(0.111)	(0.073)
Customer age	0.070***	0.116	0.008	0.043***	0.164	0.027
	(0.008)	(0.127)	(0.018)	(0.009)	(0.135)	(0.019)
Customer age²	-0.001***	-0.002	-0.000	-0.000***	-0.003	-0.000
	(0.000)	(0.003)	(0.000)	(0.000)	(0.003)	(0.000)
Customer tenure	0.002***	0.005***	-0.001	0.000	0.004***	-0.002***
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Customer tenure²	-0.000*	-0.000***	0.000*	0.000	-0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.040	0.093**	-0.012	0.062*	0.101**	0.020
	(0.033)	(0.045)	(0.042)	(0.033)	(0.044)	(0.044)
Retirement product	0.066	0.027	0.162*	0.094*	0.032	0.229**
	(0.057)	(0.076)	(0.085)	(0.057)	(0.077)	(0.094)
Home loan product	0.504***	0.730***	0.462***	0.344***	0.514***	0.323***
	(0.048)	(0.127)	(0.060)	(0.061)	(0.140)	(0.070)
Personal loan product	-0.337***	-0.459***	-0.232***	-0.223***	-0.401***	-0.073
	(0.039)	(0.060)	(0.053)	(0.042)	(0.062)	(0.056)
Savings product	0.277***	0.408***	0.178***	0.244***	0.376***	0.147***
	(0.039)	(0.060)	(0.050)	(0.040)	(0.061)	(0.051)
Term deposit product	0.180	0.382**	0.059	-0.025	0.240	-0.165
	(0.116)	(0.157)	(0.154)	(0.130)	(0.183)	(0.165)
Transaction product	0.447***	0.155	0.541***	0.987***	1.081***	0.937***
	(0.102)	(0.220)	(0.118)	(0.100)	(0.228)	(0.121)
Home insurance policy	-0.023	0.036	-0.029	-0.023	0.149	-0.066
	(0.076)	(0.160)	(0.088)	(0.085)	(0.169)	(0.098)
Motor insurance policy	0.037	0.121	0.024	0.032	0.142	0.013
	(0.106)	(0.164)	(0.124)	(0.109)	(0.168)	(0.122)
Constant	3.448***	2.704*	5.045***	3.561***	1.561	4.381***
	(0.204)	(1.560)	(0.389)	(0.239)	(1.639)	(0.422)
Observations	125,902	63,661	62,241	131,055	65,737	65,318
Number of Customers	16,487	8,243	8,244	16,487	8,243	8,244
Data treatment for closed accounts	Missing	Missing	Missing	Zero	Zero	Zero
Sample	All	Age≤28	Age>28	All	Age≤28	Age>28
R-squared	0.0316	0.0467	0.0244	0.026	0.037	0.028
Pred(Y): Non-Promotion: Control	$196.15	$161.56	$240.16	$191.83	$154.04	$240.16
Pred(Y): Non-Promotion: Treatment	$215.36	$170.07	$278.58	$213.75	$161.40	$290.41
Pred(Y): Promotion: Control	$180.53	$165.80	$197.75	$135.42	$134.09	$137.66
Pred(Y): Promotion: Treatment	$178.11	$161.69	$195.70	$133.74	$130.96	$136.63

Table A1.14: Effects of the treatment on monthly spend including customers who viewed inconsistent benefits information for a type of the low fee card. Columns (1) – (3) show the results of panel data regression models for all customers as well as younger and older customers, with spend values for months after cancellation set to missing. Columns (4) – (6) show the same specifications with spend values for months after cancellation set to zero. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the bottom section of the table. The results indicate that customers exposed to the treatment spent more on their cards, that these effects were stronger for more experienced customers, and that these effects were attenuated by the promotion, which is consistent with H3, H5, and H7.

	(1) Pr(Retain6)	(2) Pr(Retain6)	(3) Pr(Retain6)	(4) Pr(Retain9)	(5) Pr(Retain9)	(6) Pr(Retain9)
Treatment	0.281	0.084	0.508**	0.234**	0.074	0.462***
	(0.190)	(0.333)	(0.215)	(0.110)	(0.180)	(0.161)
Promotion	-1.777***	-1.718***	-1.791***	-1.313***	-0.902**	-1.566***
	(0.254)	(0.628)	(0.265)	(0.231)	(0.421)	(0.318)
Treatment x promotion	-0.269	-0.067	-0.483**	-0.227	0.102	-0.577**
	(0.208)	(0.356)	(0.242)	(0.158)	(0.245)	(0.226)
Customer age	-0.221***	-0.094	0.065	-0.100**	0.328	0.019
	(0.036)	(0.328)	(0.044)	(0.040)	(0.315)	(0.056)
Customer age²	0.003***	-0.001	-0.000	0.001***	-0.008	0.000
	(0.001)	(0.007)	(0.001)	(0.001)	(0.007)	(0.001)
Customer tenure	-0.008***	-0.010***	-0.009***	-0.005***	-0.007**	-0.006***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.003)	(0.002)
Customer tenure²	0.000***	0.000***	0.000***	0.000***	0.000*	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.048	0.028	0.059	0.075	0.040	0.111
	(0.072)	(0.109)	(0.094)	(0.089)	(0.099)	(0.136)
Retirement product	0.178	0.101	0.270	-0.106	-0.209	-0.027
	(0.141)	(0.215)	(0.208)	(0.151)	(0.178)	(0.223)
Home loan product	-0.617***	-0.754***	-0.575***	-0.550***	-0.439	-0.558***
	(0.106)	(0.189)	(0.125)	(0.170)	(0.281)	(0.193)
Personal loan product	0.758***	0.389**	1.022***	0.249**	0.001	0.493***
	(0.109)	(0.159)	(0.162)	(0.111)	(0.166)	(0.164)
Savings product	-0.201**	-0.114	-0.221**	-0.139	-0.078	-0.182
	(0.081)	(0.167)	(0.091)	(0.095)	(0.166)	(0.117)
Term deposit product	-0.569***	-0.547	-0.565***	-0.619***	-0.221	-0.776***
	(0.166)	(0.358)	(0.201)	(0.236)	(0.488)	(0.288)
Transaction product	1.581***	2.239***	1.284***	1.326***	1.852***	1.034***
	(0.125)	(0.261)	(0.161)	(0.224)	(0.432)	(0.283)
Home insurance policy	0.106	0.527	0.017	-0.046	0.736	-0.186
	(0.151)	(0.380)	(0.159)	(0.173)	(0.662)	(0.204)
Motor insurance policy	0.101	0.597	-0.029	0.031	0.738	-0.120
	(0.194)	(0.494)	(0.221)	(0.196)	(0.559)	(0.231)
Constant	7.519***	6.340	1.907**	3.945***	-0.468	1.241
	(0.683)	(4.068)	(0.893)	(0.775)	(3.809)	(1.229)
Observations	16,487	8,173	8,244	6,565	3,373	3,192
Sample	All	Age≤28	Age>28	All	Age≤28	Age>28
Pseudo R2	0.0823	0.0884	0.0965	0.0528	0.0461	0.0947
Pred(Y): Non-Promotion: Control	98.11%	98.38%	97.75%	93.34%	93.11%	93.17%
Pred(Y): Non-Promotion: Treatment	98.56%	98.50%	98.62%	94.64%	93.56%	95.53%
Pred(Y): Promotion: Control	90.33%	92.22%	88.66%	79.73%	84.93%	75.75%
Pred(Y): Promotion: Treatment	90.43%	92.34%	88.89%	79.84%	86.98%	73.78%

Table A1.15: Effects of the treatment on customer retention rates including customers who viewed inconsistent benefits information for a type of the low fee card. Effects of the treatment on customer retention rates. Columns (1) – (3) and columns (4) – (6) show the results of cross-sectional data regression models of all, younger, and older customers for retention six and nine months after card activation, respectively. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the bottom section of the table. The results indicate that transparency can lead to increased customer retention, which is consistent with H4, that the effects are strongest for more experienced customers, which is consistent with H6, and that the effects are dampened during the promotion, which is consistent with H8.

Customer Engagement Robustness Analysis: Uncontrolled Analyses

In the primary analysis presented in the manuscript, we included a series of customer-level control variables that were collected prior to the customer's exposure to the experimental manipulation. **Tables A1.16-17** demonstrate that the effects are substantively similar if these controls are excluded.

	(1) Ln(Spend)	(2) Ln(Spend)	(3) Ln(Spend)	(4) Ln(Spend)	(5) Ln(Spend)	(6) Ln(Spend)
Treatment	0.066	0.004	0.136*	0.082	0.001	0.172***
	(0.053)	(0.086)	(0.069)	(0.056)	(0.092)	(0.062)
Promotion	0.194***	0.258***	0.113	0.091	0.189**	-0.021
	(0.057)	(0.075)	(0.073)	(0.061)	(0.080)	(0.082)
Treatment x promotion	-0.076	-0.008	-0.150*	-0.094	0.001	-0.197**
	(0.062)	(0.103)	(0.084)	(0.065)	(0.110)	(0.080)
Constant	5.096***	4.927***	5.278***	4.968***	4.815***	5.133***
	(0.047)	(0.063)	(0.059)	(0.047)	(0.066)	(0.061)
Observations	121,714	61,231	60,483	126,696	63,237	63,459
Customers	15,947	7,932	8,015	15,947	7,932	8,015
Data treatment for closed accounts	Missing	Missing	Missing	Zero	Zero	Zero
Sample	All	Age≤28	Age>28	All	Age≤28	Age>28
R-squared	0.0015	0.0038	0.0004	0.0002	0.0018	0.001
Pred(Y): Non-Promotion: Control	$163.37	$137.94	$195.96	$143.78	$123.32	$169.51
Pred(Y): Non-Promotion: Treatment	$174.48	$138.49	$224.51	$156.04	$123.49	$201.42
Pred(Y): Promotion: Control	$198.32	$178.47	$219.33	$157.45	$149.01	$165.92
Pred(Y): Promotion: Treatment	$196.35	$177.81	$216.18	$155.62	$149.32	$161.94

Table A1.16: Effects of the treatment on monthly spend excluding control variables. Columns (1) – (3) show the results of panel data regression models for all customers as well as younger and older customers, with spend values for months after cancellation set to missing. Columns (4) – (6) show the same specifications with spend values for months after cancellation set to zero. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the bottom section of the table. The results indicate that customers exposed to the treatment spent more on their cards, that these effects were stronger for more experienced customers, and that these effects were attenuated by the promotion, which is consistent with H3, H5, and H7.

	(1) Pr(Retain6)	(2) Pr(Retain6)	(3) Pr(Retain6)	(4) Pr(Retain9)	(5) Pr(Retain9)	(6) Pr(Retain9)
Treatment	0.267	0.094	0.451*	0.248**	0.103	0.416***
	(0.191)	(0.320)	(0.232)	(0.114)	(0.183)	(0.157)
Promotion	-1.009***	-0.846***	-1.116***	-0.713***	-0.670***	-0.751***
	(0.128)	(0.198)	(0.174)	(0.123)	(0.153)	(0.192)
Treatment x promotion	-0.276	-0.084	-0.475*	-0.251	0.142	-0.594***
	(0.209)	(0.344)	(0.254)	(0.166)	(0.251)	(0.227)
Constant	3.485***	3.586***	3.385***	2.482***	2.525***	2.439***
	(0.114)	(0.177)	(0.160)	(0.080)	(0.120)	(0.128)
Observations	15,947	7,932	8,015	6,316	3,231	3,085
Customers	All	Age≤28	Age>28	All	Age≤28	Age>28
Pseudo R2	0.0221	0.0144	0.0284	0.0241	0.0129	0.0395
Pred(Y): Non-Promotion: Control	97.02%	97.30%	96.72%	92.29%	92.59%	91.97%
Pred(Y): Non-Promotion: Treatment	97.71%	97.54%	97.89%	93.88%	93.26%	94.55%
Pred(Y): Promotion: Control	92.24%	93.93%	90.63%	85.44%	86.47%	84.39%
Pred(Y): Promotion: Treatment	92.18%	93.99%	90.43%	85.41%	89.08%	81.90%

Table A1.17: Effects of the treatment on customer retention rates excluding control variables. Effects of the treatment on customer retention rates. Columns (1) – (3) and columns (4) – (6) show the results of cross-sectional data regression models of all, younger, and older customers for retention six and nine months after card activation, respectively. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively. Marginal effect estimates are provided for each condition in the bottom section of the table. The results indicate that transparency can lead to increased customer retention, which is consistent with H4, that the effects are strongest for more experienced customers, which is consistent with H6, and that the effects are dampened during the promotion, which is consistent with H8.

Customer Engagement Robustness Analysis: Truncated Spend Analysis

In the primary spend analysis, presented in the manuscript, we present panel analyses based on the full sample of data. In **Table A1.18**, below, we present the same analysis, regressed on the first six months of observations for each customer. This approach enables us to analyze a sample of equivalent length from every customer in our study. We note that the results are substantively similar to the results presented in the primary analysis. The sole difference is that with the diminished power afforded by having fewer observations, we see an insignificant main effect of spend on the full sample, but the effect emerges for the high-experience subgroup analysis.

	(1) Ln(Spend)	(2) Ln(Spend)	(3) Ln(Spend)	(4) Ln(Spend)	(5) Ln(Spend)	(6) Ln(Spend)
Treatment	0.071	0.035	0.120*	0.076	0.030	0.137**
	(0.050)	(0.073)	(0.067)	(0.052)	(0.077)	(0.061)
Promotion	-0.159	-0.005	-0.316*	-0.361***	-0.132	-0.592***
	(0.115)	(0.132)	(0.167)	(0.122)	(0.126)	(0.177)
Treatment x promotion	-0.088	-0.040	-0.158*	-0.093	-0.032	-0.176**
	(0.059)	(0.093)	(0.081)	(0.061)	(0.097)	(0.078)
Customer age	0.071***	0.099	0.013	0.050***	0.136	0.027
	(0.008)	(0.117)	(0.017)	(0.009)	(0.124)	(0.018)
Customer age²	-0.001***	-0.001	-0.000	-0.001***	-0.002	-0.000*
	(0.000)	(0.002)	(0.000)	(0.000)	(0.003)	(0.000)
Customer tenure	0.002***	0.005***	-0.001	0.000	0.004***	-0.002***
	(0.000)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)
Customer tenure²	-0.000	-0.000***	0.000*	0.000	-0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.070**	0.112***	0.027	0.093***	0.115***	0.067*
	(0.032)	(0.043)	(0.040)	(0.032)	(0.042)	(0.041)
Retirement product	0.082	0.033	0.195**	0.120**	0.049	0.273***
	(0.056)	(0.072)	(0.084)	(0.056)	(0.073)	(0.090)
Home loan product	0.475***	0.715***	0.432***	0.343***	0.514***	0.323***
	(0.048)	(0.124)	(0.059)	(0.059)	(0.136)	(0.067)
Personal loan product	-0.221***	-0.341***	-0.123**	-0.120***	-0.279***	0.010
	(0.038)	(0.062)	(0.049)	(0.039)	(0.062)	(0.052)
Savings product	0.249***	0.346***	0.173***	0.226***	0.318***	0.155***
	(0.039)	(0.058)	(0.048)	(0.039)	(0.059)	(0.048)
Term deposit product	0.137	0.320**	0.031	-0.049	0.189	-0.172
	(0.114)	(0.145)	(0.150)	(0.125)	(0.165)	(0.160)
Transaction product	0.503***	0.249	0.586***	0.946***	0.954***	0.928***
	(0.100)	(0.218)	(0.115)	(0.095)	(0.214)	(0.116)
Home insurance policy	-0.057	0.041	-0.076	-0.067	0.124	-0.116
	(0.068)	(0.150)	(0.080)	(0.077)	(0.158)	(0.088)
Motor insurance policy	0.072	0.152	0.066	0.054	0.147	0.046
	(0.097)	(0.153)	(0.117)	(0.101)	(0.159)	(0.117)
Constant	3.540***	2.923**	5.099***	3.592***	2.062	4.530***
	(0.214)	(1.449)	(0.398)	(0.234)	(1.507)	(0.415)
Observations	90,208	45,165	45,043	92,696	46,098	46,598
Data treatment for closed accounts	Missing	Missing	Missing	Zero	Zero	Zero
Customers	15,942	7,932	8,010	15,942	7,932	8,010

Table A1.18: Spend analysis, restricted to the first six months of usage. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. Columns (1) – (3) show the results of panel data regression models for all customers as well as younger and older customers, with spend values for months after cancellation set to missing. Columns (4) – (6) show the same specifications with spend values for months after cancellation set to zero. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively.

Testing Experience as a Potential Moderator of the Effects of Tradeoff Transparency

In the primary analyses presented in the manuscript, we demonstrate positive main effects of tradeoff transparency on customer spending and retention, and we note that the effects are particularly pronounced among more experienced customers. However, a different, yet related, question is whether experience moderates the relationships between tradeoff transparency and these measures of customer engagement.

To test these relationships, we create a moderator variable, called $OVER28_i$, which denotes a more experienced customer. As noted earlier in the Appendix, the probability a 28-year old prospective customer had a credit card at the time of this study was 58.5%. Prospective customers who were older than 28 years old were more likely to have a credit card, and prospective customers who were younger were less likely to have one.

In the first specification presented below, γ_5 captures the degree to which tradeoff transparency has a differential effect on the spending of more experienced customers during the non-promotion period:

$$\ln\left(SPEND_{it}\right) = f\begin{pmatrix} \gamma_0 + \gamma_1 TREAT_i + \gamma_2 PROMO_i + \gamma_3 TREAT_i \times PROMO_i + \gamma_4 OVER28_i + \\ \gamma_5 TREAT_i \times OVER28_i + \gamma_6 PROMO_i \times OVER28_i + \\ \gamma_7 TREAT_i \times OVER28_i \times PROMO_i + \gamma_8 AGE_i + \gamma_9 AGE_i^2 + \gamma_{10} TENURE_i + \\ \gamma_{11} TENURE_i^2 + \gamma_{12} GENDER_i + \gamma_{13} HLOAN_i + \gamma_{14} PLOAN_i + \gamma_{15} TRANS_i + \\ \gamma_{16} SAV_i + \gamma_{17} HINS_i + \gamma_{18} VINS_i + \gamma_{19} RET_i + X_i + \epsilon_{it} \end{pmatrix} \tag{4}$$

In the next specification presented below, δ_5 captures the degree to which tradeoff transparency has a differential effect on the retention of more experienced customers during the non-promotion period:

$$\Pr\left(RETAIN_i\right) = f\begin{pmatrix} \delta_0 + \delta_1 TREAT_i + \delta_2 PROMO_i + \delta_3 TREAT_i \times PROMO_i + \delta_4 OVER28_i + \\ \delta_5 TREAT_i \times OVER28_i + \delta_6 PROMO_i \times OVER28_i + \\ \delta_7 TREAT_i \times OVER28_i \times PROMO_i + \delta_8 AGE_i + \delta_9 AGE_i^2 + \delta_{10} TENURE_i + \\ \delta_{11} TENURE_i^2 + \delta_{12} GENDER_i + \delta_{13} HLOAN_i + \delta_{14} PLOAN_i + \delta_{15} TRANS_i + \\ \delta_{16} SAV_i + \delta_{17} HINS_i + \delta_{18} VINS_i + \delta_{19} RET_i + X_i + \epsilon_i \end{pmatrix} \tag{5}$$

We present the results of these analyses in **Table A1.19.** The results demonstrate that experience does not moderate the relationship between tradeoff transparency and spending, neither when spending is set to missing during months following cancellation (γ=0.117, p=0.31), or when spending is set to 0 during months following cancellation (γ=0.158, p=0.16). Similarly, the results reveal that experience does not moderate the relationship between tradeoff transparency, neither after six months (δ=0.324, p=0.44), nor after nine months (δ=0.328, p=0.21).

	(1) Ln(Spend)	(2) Ln(Spend)	(3) Pr(Retain6)	(4) Pr(Retain9)
Treatment	0.025	0.020	0.105	0.096
	(0.085)	(0.092)	(0.327)	(0.180)
Promo	-0.035	-0.271**	-1.573***	-1.301***
	(0.114)	(0.124)	(0.295)	(0.252)
Treatment x Promo	-0.034	-0.021	-0.065	0.165
	(0.103)	(0.111)	(0.352)	(0.254)
Over 28	0.222**	0.196**	0.054	-0.109
	(0.087)	(0.094)	(0.274)	(0.229)
Treatment x Over 28	0.117	0.158	0.324	0.328
	(0.115)	(0.113)	(0.422)	(0.262)
Promo x Over 28	-0.138	-0.197*	-0.271	-0.090
	(0.095)	(0.103)	(0.269)	(0.253)
Treatment x Promo x Over 28	-0.134	-0.186	-0.353	-0.734**
	(0.142)	(0.144)	(0.451)	(0.354)
Customer age	0.053***	0.036***	-0.186***	-0.065
	(0.012)	(0.012)	(0.050)	(0.045)
Customer age²	-0.001***	-0.000***	0.003***	0.001*
	(0.000)	(0.000)	(0.001)	(0.001)
Customer tenure	0.002***	0.000	-0.009***	-0.005***
	(0.000)	(0.001)	(0.001)	(0.002)
Customer tenure²	-0.000	0.000	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
Male indicator	0.052	0.073**	0.045	0.075
	(0.034)	(0.034)	(0.074)	(0.090)
Retirement product	0.067	0.104*	0.241	-0.121
	(0.058)	(0.058)	(0.155)	(0.154)
Home loan product	0.509***	0.349***	-0.595***	-0.565***
	(0.049)	(0.062)	(0.105)	(0.166)
Personal loan product	-0.324***	-0.209***	0.767***	0.240**
	(0.040)	(0.043)	(0.110)	(0.105)
Savings product	0.279***	0.249***	-0.187**	-0.124
	(0.040)	(0.040)	(0.080)	(0.096)
Term deposit product	0.192	-0.029	-0.614***	-0.650***
	(0.119)	(0.132)	(0.167)	(0.230)
Transaction product	0.496***	1.041***	1.606***	1.402***
	(0.105)	(0.101)	(0.126)	(0.225)
Home insurance policy	-0.039	-0.046	0.080	-0.084
	(0.076)	(0.085)	(0.154)	(0.169)
Motor insurance policy	0.043	0.036	0.078	0.029
	(0.105)	(0.108)	(0.197)	(0.199)
Constant	3.625***	3.554***	6.810***	3.238***
	(0.250)	(0.265)	(0.911)	(0.818)
Observations	121,679	126,658	15,942	6,314
Customers	15,942	15,942	15,942	6,314
Data treatment for closed accounts	Missing	0	NA	NA
R-squared (Pseudo R-squared)	0.0174	0.0143	0.0854	0.0580

Table A1.19: Experience does not moderate the relationship between tradeoff transparency and spending, nor between tradeoff transparency and retention. All models include indicator variables for the week the credit card was activated, and are estimated with robust standard errors clustered by activation date. *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively.

Testing Trust as a Potential Mechanism Underlying the Effects of Tradeoff Transparency

We recruited 201 participants (62.2% male, *M*=39.18 years old) on the Amazon Mechanical Turk platform. After completing the informed consent process, participants were shown the credit card detail pages from two competing banks – Purple Bank and Yellow Bank – both of which offered a "Low Fee Credit Card." Participants were instructed to carefully review the two webpages, as if they were considering applying for a credit card from these two banks.

The details of the credit card offerings were identical to the Low Fee Credit Card offered by CBA in our field experiment. For each participant, we randomly assigned which of the two hypothetical banks offered tradeoff transparency, and we counterbalanced the order of each bank's presentation (**Figure A1.4**). After reviewing the two banks' webpages, every participant was asked a pair of forced choice questions – "Which bank is more trustworthy?" and "If you were given the opportunity, from which bank would you be more likely to apply for a credit card?"

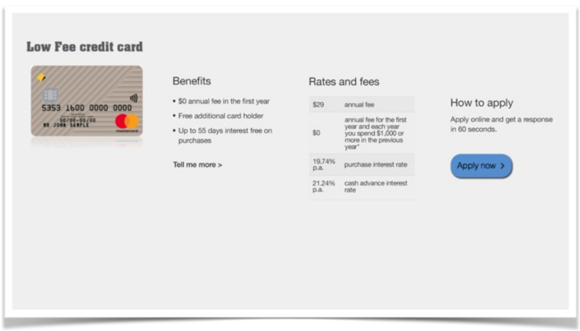

Yellow Bank – Control (benefits only)

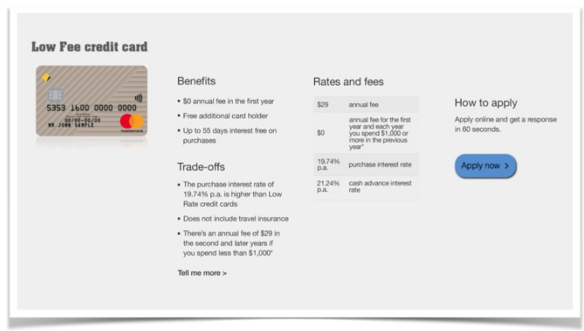

Yellow Bank – Tradeoff transparency

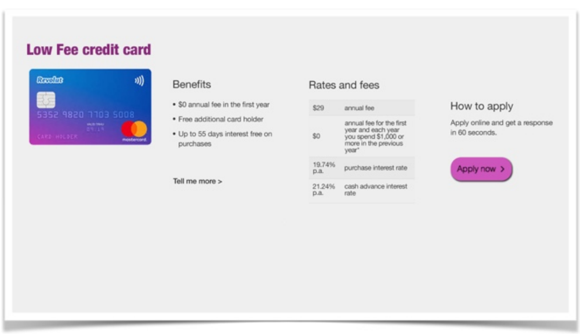

Purple Bank – Control (benefits only)

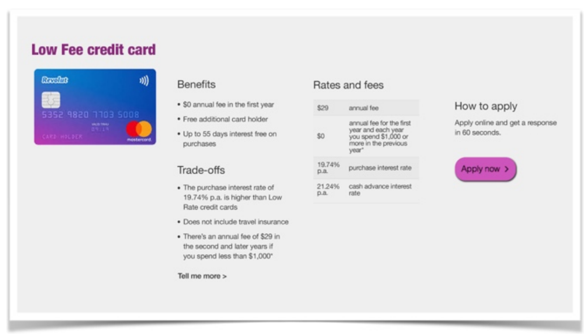

Purple Bank – Tradeoff transparency

Figure A1.4: Screenshots of Credit Card Websites Presented in the Online Experiment (n=201).

We test the impact of tradeoff transparency on trust and willingness to apply for a credit card using one-sided binomial tests of the null hypothesis that tradeoff transparency has no effect. If participant responses were the product of random chance, we would expect to see participants choose the bank leveraging tradeoff transparency as being more trustworthy and being their preferred bank half of the time. However, 71% of participants reported that the bank that exhibited tradeoff transparency was more trustworthy ($p<0.01$, one-sided). Moreover, 69% reported a higher intention to apply for a credit card from the bank that provided tradeoff transparency ($p<0.01$, one-sided). Finally, consistent with prior research, there was a high correlation between trust and preference, such that the bank that was identified as being more trustworthy was also selected disproportionately as the participant's choice as the bank from which they would be more likely to apply for a credit card (ρ =0.794, $p<0.01$) (**Figure A1.5**).

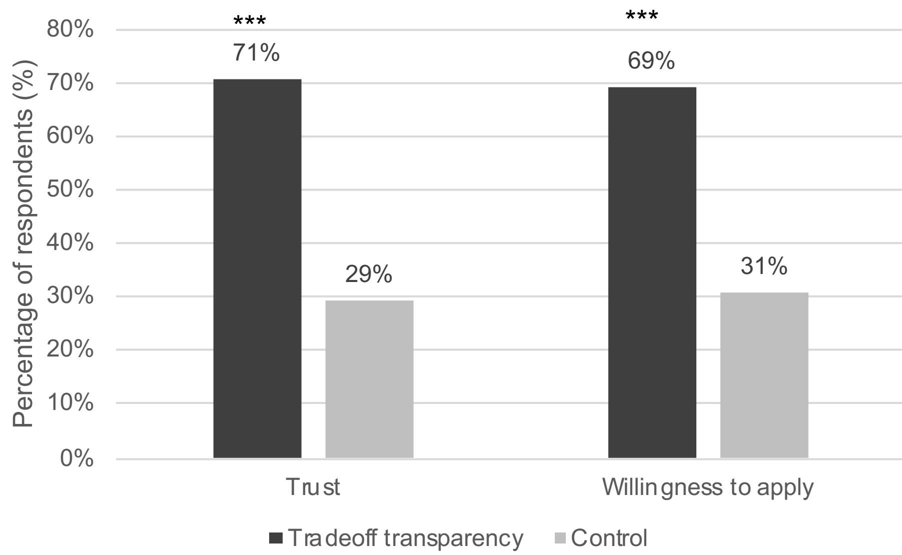

Figure A1.5: Effects of Tradeoff Transparency on Trust and Willingness to Apply for a Credit Card (n=201). *, **, and *** signify significance at the 10%, 5%, and 1% levels respectively, drawn from one-sided binomial tests of the null hypothesis that tradeoff transparency has no effect on trust and willingness to apply for a credit card. The results provide evidence that tradeoff transparency can engender trust and brand preference in customers. Moreover, consistent with prior literature, we note a positive and statistically significant correlation between trust and willingness to apply (ρ=0.794, p<0.01).

Additional Screenshots from the Treatment and Control Conditions

Figures A1.6-A1.7 present additional screenshots from the treatment and control condition of the primary experiment. The experimental manipulation involved more than 50 blocks of content on more than 20 pages of CBA's public-facing and secure online banking websites, such that every on credit card marketing webpage where the features and benefits of a credit card were described, so too were its tradeoffs for customers randomly assigned to the treatment. The nature of the tradeoffs manipulation varied from page template to page template, to complement the way benefits were presented on various parts of the website.

Control Condition

Only the benefits of each credit card are presented on this page

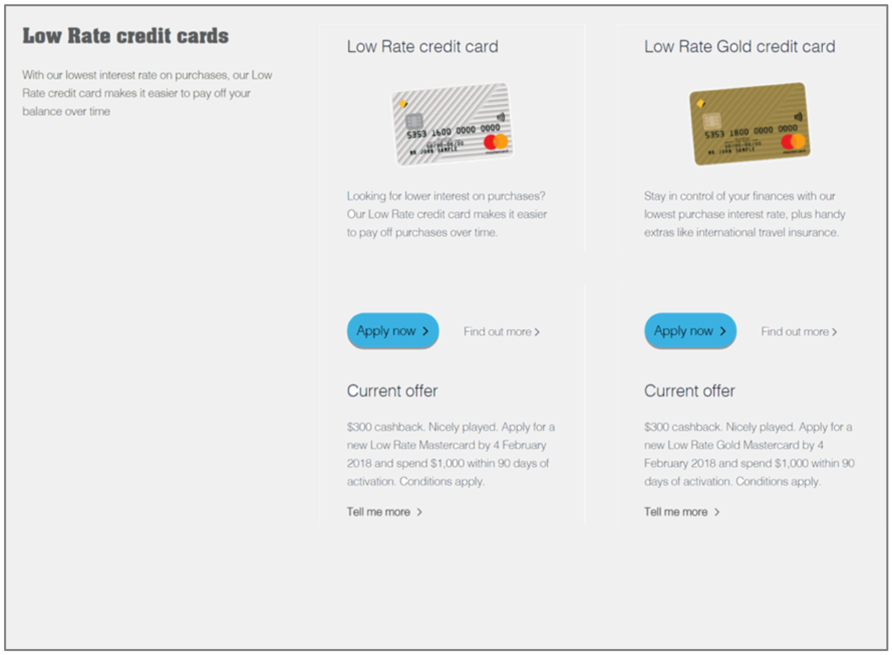

Treatment Condition

The benefits and the tradeoffs, in terms of the annual fees, are presented

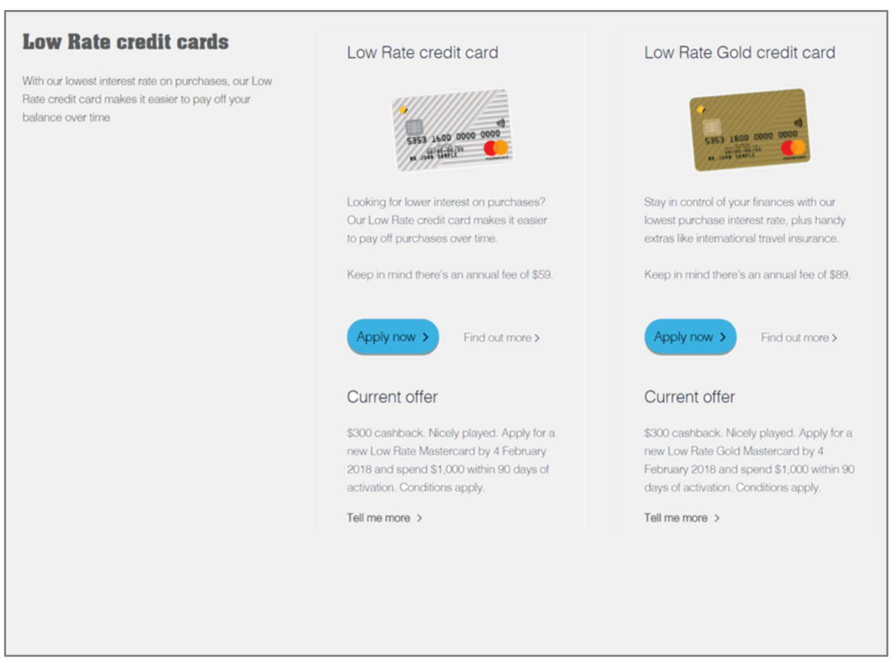

Figure A1.6. Additional example credit card marketing pages in the control and treatment conditions for a page presenting the two cards in the low rate credit card family.

Control Condition

Only the benefits of each credit card are presented in the product detail pages

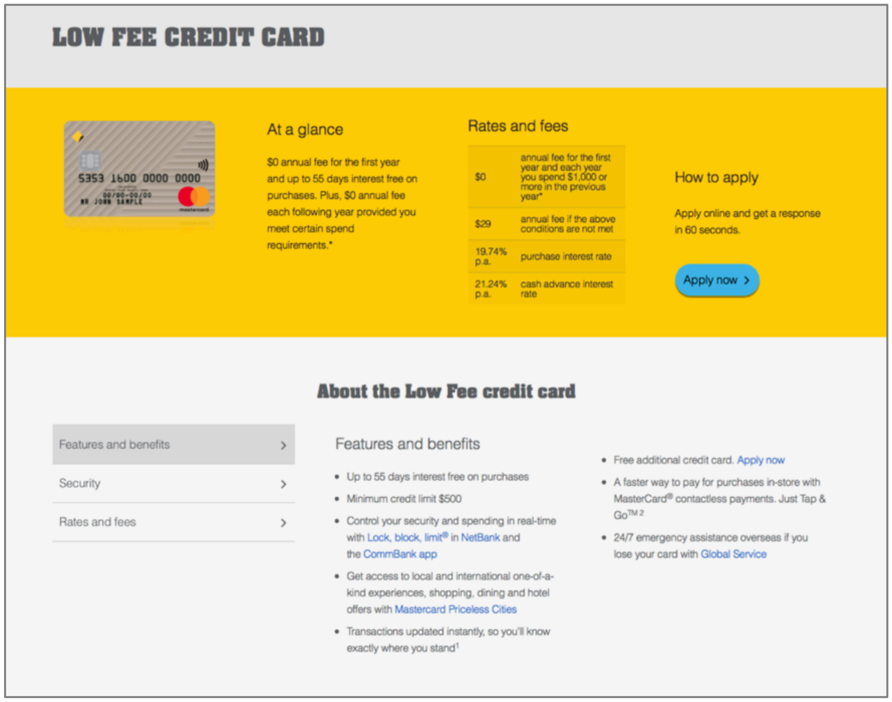

Treatment Condition

In the treatment condition, a section was added beneath the Features and benefits tab, which presents the card's tradeoffs

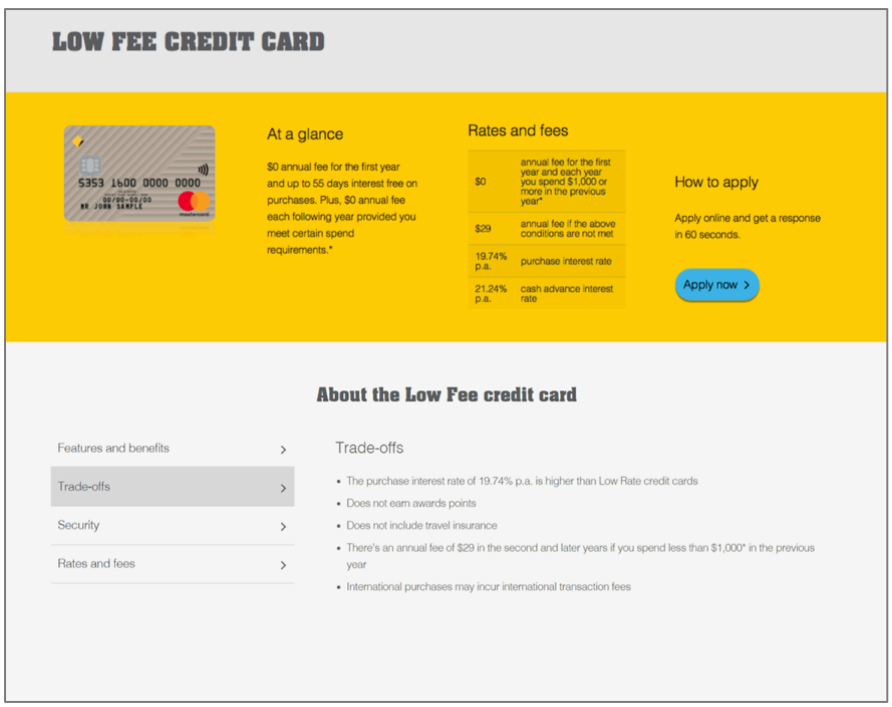

Figure A1.7. Additional example credit card product detail pages in the control and treatment conditions for a credit card detail page. For the low fee credit card detail page, in the treatment condition, an entire section was added to the left panel, including the tradeoffs, in addition to the features and benefits.

Online Appendix

Table A2.1. Summary Statistics for Delayed Orders. All of the duration statistics are in hours.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Duration	28,574	41.174	28.022	14.773	24.843	32.018	46.747	528.762
First leg duration	28,574	16.621	24.976	0.000	2.067	9.038	20.006	515.219
Second leg duration	28,574	13.471	9.789	1	8	13	17	231
Third leg duration	28,574	11.082	16.180	−0	3	5	12	319
No. hours delayed	28,574	21.706	27.380	515	24	13.0	4	0
Duration/Time Allotted	28,574	2.151	1.450	1.001	1.220	1.801	2.277	40.009

Table A2.2. Summary Statistics for On-Time Orders

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Duration	98,480	16.095	5.349	2.051	12.352	16.795	20.256	27.988
First leg duration	98,480	2.072	2.656	0.000	0.446	0.965	2.285	24.282
Second leg duration	98,480	10.044	4.488	0	6	10	14	25
Third leg duration	98,480	3.979	3.071	0	2	3	5	23
No. hours early	98,480	4.845	2.302	0	3	5	6	25
Duration/Time Allotted	98,480	0.750	0.145	0.081	0.676	0.779	0.853	1.000

Table A2.3. The 95th percentile first leg durations for all on-time orders, separated by the hour of the day in which they were placed. N indicates the total number of observations across March 1 to March 31. If an order takes longer than the value for the corresponding hour when it was placed, it is considered late in the first leg.

Order Hour	First Leg Duration (95 Pct)	N
0	9.93	3569
1	8.99	1424
2	8.18	571
3	7.15	311
4	5.98	269
5	4.87	346
6	3.64	1079
7	2.95	2548
8	2.29	4455
9	1.59	5685
10	0.92	5588
11	6.89	6106
12	5.75	6317
13	5.30	6615
14	4.91	6333
15	4.61	5904
16	3.97	5355
17	3.47	4343
18	2.87	3890
19	2.27	4249
20	1.87	5050
21	1.50	5800
22	0.93	6004
23	10.93	6669

Table A2.4. Summary Statistics for Order Features (N=127,054)

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Total order cost	115.218	187.458	−30	50	149	28,912
No. distinct SKUs	1.119	0.386	1	1	1	8
No. items in order	1.358	1.766	1	1	1	170
Contains gift	0.068	0.251	0	0	0	1
Contains direct discount	0.748	0.434	0	0	1	1
Contains quantity discount	0.294	0.455	0	0	1	1
Contains bundle discount	0.020	0.141	0	0	0	1
Used a coupon	0.250	0.433	0	0	1	1

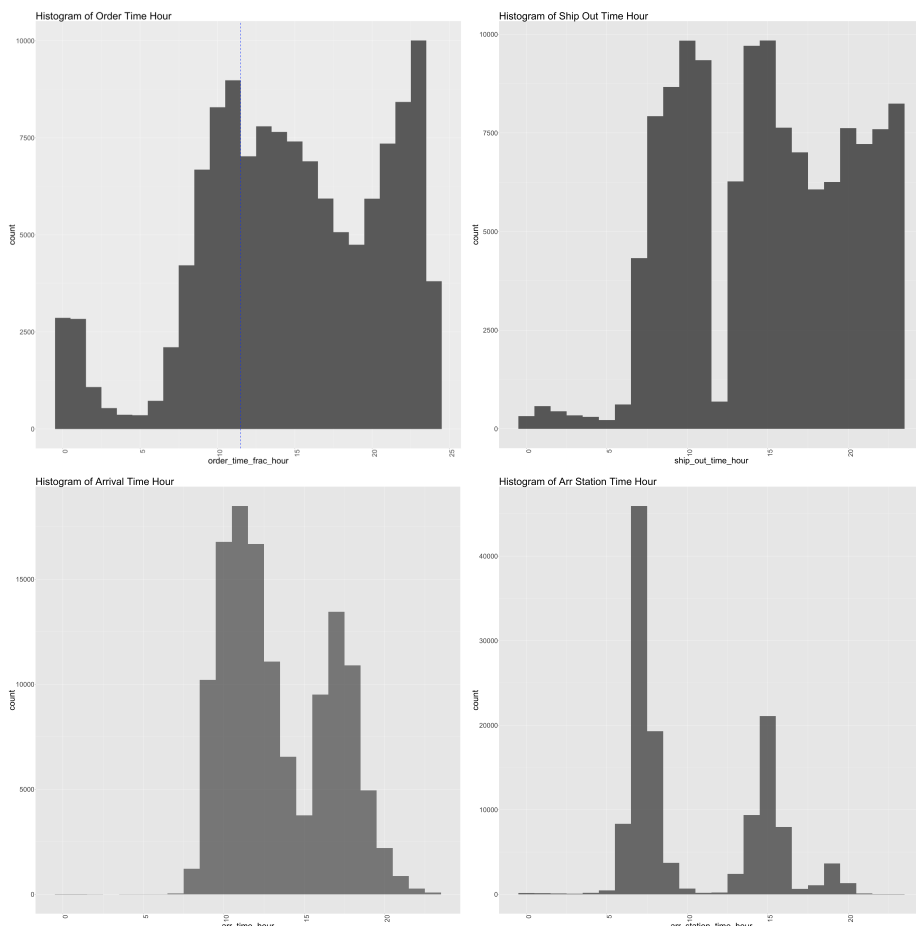

Figure A2.1. (clockwise from top left) Histograms of order times, ship out times, arrival station times, and arrival times. X-axis represents the hour (from 0:00-23:00), and the y-axis represents the frequency.

Table A2.5. Robustness check: RD model for duration of each leg for orders placed between 10 am and 12 pm, with promise of same- and next-day delivery, including full list of control variables. The coefficients for order cost, user education level, number of gifts, and user age were not significant variables in predicting duration. Robustness checks for other outcome variables (e.g., $P(\text{Delay})$ and $P(\text{Late leg}_t)$) show similar patterns as without the full list of control variables.

	Dependent variable:					
	D_1^{Ontime}	D_2^{Ontime}	D_3^{Ontime}	D_1^{Delay}	D_2^{Delay}	D_3^{Delay}
	(1)	(2)	(3)	(4)	(5)	(6)
AFTER_11	3.10^{***} (0.03)	10.38^{***} (0.13)	2.95^{***} (0.10)	19.54^{***} (1.72)	8.66^{***} (0.83)	14.22^{***} (1.96)
τ	-1.00^{***} (0.05)	−0.07 (0.17)	−0.17 (0.13)	0.30 (1.30)	0.34 (0.39)	−1.89 (1.58)
AFTER_11$\cdot\tau$	−0.05 (0.06)	−0.27 (0.24)	0.62^{***} (0.19)	-6.31^{**} (3.03)	−0.44 (1.40)	3.80 (3.34)
(type)1	-0.19^{***} (0.04)	-0.30^{*} (0.16)	0.35^{***} (0.12)	3.61^{**} (1.62)	-7.78^{***} (0.44)	−1.77 (1.73)
num_items	−0.002 (0.01)	−0.0002 (0.03)	0.03 (0.02)	0.07 (0.27)	-0.20^{**} (0.10)	0.29 (0.50)
num_skus	0.02 (0.03)	0.15 (0.10)	0.03 (0.08)	-1.75^{*} (1.00)	0.45 (0.33)	−0.74 (1.12)
(city_level)1	-0.26^{***} (0.07)	-0.64^{**} (0.29)	1.12^{***} (0.22)	−1.11 (2.84)	-1.51^{*} (0.84)	2.60 (3.02)
(city_level)2	−0.09 (0.07)	−0.32 (0.29)	0.38^{*} (0.22)	−0.43 (2.83)	−1.20 (0.84)	5.90^{*} (3.02)
(city_level)3	-0.17^{**} (0.08)	0.69^{**} (0.30)	−0.01 (0.23)	0.88 (2.97)	−0.46 (0.88)	4.84 (3.33)
(city_level)4	−0.12 (0.08)	0.34 (0.32)	0.07 (0.25)	−1.91 (3.18)	0.14 (0.95)	2.93 (3.54)
(city_level)5	0.53^{***} (0.18)	−0.25 (0.68)	−0.14 (0.54)	−0.27 (6.98)	−2.00 (2.31)	2.92 (9.10)
contains_gift	−0.04 (0.11)	−0.41 (0.39)	0.43 (0.32)	13.99^{***} (4.06)	−0.01 (1.56)	0.87 (4.72)
direct_discount	−0.01 (0.02)	0.07 (0.08)	0.03 (0.07)	1.46 (0.91)	0.06 (0.28)	1.00 (1.00)
quantity_discount	0.08^{***} (0.02)	-0.18^{**} (0.08)	0.16^{***} (0.06)	0.91 (0.81)	0.10 (0.25)	−0.04 (0.91)
bundle_discount	-0.12^{**} (0.06)	−0.04 (0.23)	0.08 (0.19)	2.29 (2.67)	−0.28 (0.88)	1.93 (2.80)
used_coupon	0.02 (0.02)	−0.04 (0.08)	−0.05 (0.06)	0.16 (0.83)	-0.42^{*} (0.26)	0.12 (0.91)
(user_level)0	0.22 (0.47)	3.05 (1.88)	0.85 (1.81)			−14.95 (17.61)
(user_level)1	0.03 (0.09)	1.28^{***} (0.34)	-0.78^{***} (0.27)	−3.12 (3.82)	0.37 (1.13)	2.50 (3.72)
(user_level)2	0.01 (0.09)	1.22^{***} (0.34)	-0.88^{***} (0.27)	−2.23 (3.78)	0.40 (1.12)	2.85 (3.67)
(user_level)3	0.002 (0.09)	1.14^{***} (0.34)	-0.78^{***} (0.27)	−1.48 (3.80)	0.54 (1.12)	2.61 (3.71)
(user_level)4	0.003 (0.09)	1.12^{***} (0.35)	-0.78^{***} (0.28)	−1.29 (3.87)	0.92 (1.14)	2.97 (3.78)
(user_level)10	−0.07 (0.10)	−0.15 (0.38)	−0.33 (0.31)	36.03^{***} (4.11)	0.11 (1.27)	4.60 (4.47)
(plus)1	0.003 (0.02)	−0.07 (0.09)	−0.001 (0.07)	1.24 (0.96)	0.01 (0.29)	0.17 (1.05)
(gender)M	−0.02 (0.02)	0.07 (0.08)	−0.09 (0.06)	−0.94 (0.92)	-0.49^{*} (0.28)	−0.54 (0.99)
(gender)U	0.81^{**} (0.38)	0.25 (1.30)	−0.99 (1.03)	−4.32 (11.66)	−3.76 (4.59)	−12.36 (17.26)
(married)S	−0.01 (0.02)	-0.22^{***} (0.08)	0.07 (0.06)	−0.05 (0.88)	0.04 (0.27)	0.94 (0.95)
(married)U	0.17^{*} (0.10)	-0.75^{*} (0.39)	0.53^{*} (0.30)	−1.93 (4.19)	−0.93 (1.30)	2.22 (4.92)
(pur_power)1	0.19 (0.11)	−0.64 (0.44)	0.69^{*} (0.35)	−5.05 (5.06)	−0.47 (1.56)	−6.09 (5.67)
(pur_power)2	0.23^{**} (0.09)	−0.51 (0.36)	0.30 (0.28)	−4.66 (3.89)	−1.18 (1.22)	−4.03 (4.60)
(pur_power)3	0.20^{**} (0.09)	−0.49 (0.36)	0.29 (0.28)	−4.94 (3.90)	−0.72 (1.23)	−0.94 (4.62)
(pur_power)4	0.27^{**} (0.11)	−0.25 (0.44)	0.39 (0.34)	−6.77 (4.81)	−0.36 (1.46)	−3.62 (5.51)
Observations	12,383	13,037	14,689	4,327	3,673	2,021
R^2	0.60	0.66	0.22	0.17	0.22	0.14
Adj. R^2	0.59	0.66	0.22	0.16	0.21	0.12
Resid. SE	0.91	3.66	3.08	23.01	6.45	17.01
F Statistic	363.87^{***}	507.64^{***}	82.84^{***}	18.64^{***}	21.04^{***}	6.93^{***}

Appendix A.3.1 - Process for constructing the dataset.

1. We drop all the duplicate rows in the Orders table. Given that each row should contain a unique order-SKU pair, we assume that duplicate rows contain redundant information and can be safely removed.

2. We group the Orders table by unique order_ID and aggregate the information across the SKUs in each order. This allows to get the total cost for each order and the number of items and SKUs in the order. This gives us a version of the Orders table with a single row per order.

3. On the Delivery table, we select a single row per order_ID corresponding to the last delivered package for that order. This gives us a version of the Delivery table with a single row per unique order.

4. We then merge our modified Orders and Delivery tables, giving us a single table with one row per order containing the information of which SKUs are in the order, and when it's latest package was shipped out and delivered.

5. We further filter this table for consistency on timestamps. We only keep orders in which: (i) the timestamp when the order was placed occurred before the recorded timestamp of when the order's latest package was shipped out, (ii) the ship out time occurred before the order's package arrived at the delivery station, and (iii) the package arrived at the delivery station prior to being marked as delivered. We assume that timestamps that contradict this pattern are due to faulty data capturing and therefore should be removed from our dataset.

6. We filter our resulting table to just select the rows where the order has a promise of 1 (promise of same and next day deliveries). Out of 307,009 unique orders that have a listed promise, 140,433 (46%) of them have a promise of 1.

7. We merge this table containing detailed order and delivery information with the User table, so that we have a single table with each row representing a unique order. Each row contains information on the delivery process for that order as well as the demographic data of the user who placed the order. Note that this step drops the orders for which we don't have any User information.

Table A2.6. Full list of all the features included in the Machine Learning model, listed in the order of variable importance (as measured by Gini impurity).

Full List of Features (and Corresponding Gini Impurity Values)					
Baseline		1st leg info.		1st + 2nd leg info.	
order_time_day	872.68	time_left_after _first_leg	6995.82	time_left_after_ second_leg	6555.10
time_allotted	546.66	first_leg_length	3513.60	time_left_after_ first_leg	4605.76
dc_ori	293.06	ship_out_time_ frac_hour	589.04	first_leg_length	2056.99
order_time_ frac_hour	288.14	time_allotted	391.39	second_leg_length	847.88
dc_des	264.12	order_time_day	381.30	order_time_day	198.73
order_time_ day_of_week	215.16	order_time_ frac_hour	295.42	time_allotted	191.73
total_order_cost	212.23	dc_ori	95.78	ship_out_time_ frac_hour	179.05
city_level	77.62	after_11	95.12	arr_station_time	160.20
type	68.71	dc_des	92.34	order_time_ frac_hour	140.38
after_11	48.77	total_order_cost	75.02	dc_ori	57.00
user_level	40.50	order_time_day_ of_week	58.12	dc_des	51.71
contains_quantity_ discount	39.27	city_level	33.74	total_order_cost	39.89
num_gifts	32.02	age	21.14	city_level	28.26
age	31.37	type	20.98	order_time_day_ of_week	24.45
contains_gift	31.00	user_level	18.18	after_11	20.26
purchase_power	21.27	education	11.724	age	15.289
education	21.11	contains_quantity_ discount	10.806	user_level	12.651
number_of_items	20.57	purchase_power	10.677	education	7.724
used_a_coupon	19.0	num_gifts	8.681	purchase_power	7.636
contains_direct_ discount	13.9	number_of_items	8.298	contains_quantity_ discount	7.421
number_distinct_ skus	12.6	contains_gift	7.214	type	6.736
marital_status	11.2	marital_status	6.781	marital_status	6.466
gender	9.8	used_a_coupon	6.158	number_of_items	5.869
plus	6.3	number_distinct_ skus	4.821	gender	3.766
contains_bundle_ discount	2.0	gender	4.731	num_gifts	3.753
num_pack	0.2	contains_direct_ discount	4.537	contains_gift	3.708
		plus	2.999	used_a_coupon	3.312

Appendix A: Tables and Figures

A.3.1. Correlation Tables

Table 3.5. Correlations Study 1.

	1	2
1. Willingness to sign	1	
2. Perceived value	0.698^{**}	1

Note. $N = 211$, ** $p < 0.01$

We ran Pearson product-moment correlations to determine the relationship between willingness to sign contracts, and perceived value. There was a strong and positive correlation between perceived value and willingness to sign contracts ($r = .698, N = 211, p < 0.01$), see Table 3.5.

Table 3.6. Correlations Study 2.

	1	2	3	4	5
1. Understanding	1				
2. Perc. understanding	0.134^{*}	1			
3. Trust	0.080	0.758^{**}	1		
4. Willingness to sign	0.134^{*}	0.676^{**}	0.805^{**}	1	
5. Perceived value	0.023	0.647^{**}	0.788^{**}	0.757^{**}	1

Note. $N = 222$, ** $p < 0.01$; * $p < 0.05$

Correlations Study 2. We ran Pearson product-moment correlations to determine the relationship between contract understanding, perceived contract understanding, trust, willingness to sign contracts and perceived value. All correlations were positive. There was a moderately strong correlation between contract understanding and perceived contract understanding ($r = 0.134, p < 0.05$), as well as between contract understanding and willingness to sign contracts ($r = 0.134, p < 0.05$). There was no correlation between contract understanding and trust in the company, and there was also no correlation between contract understanding and perceived value. Perceived contract understanding was highly correlated with trust in the company ($r = 0.758, p < 0.01$), willingness to sign contracts ($r = 0.676, p < 0.01$) and value perception ($r = 0.647, p < 0.01$). Moreover, trust in the company highly correlated with willingness to sign contracts ($r = 0.805, p < 0.01$), and perceived value ($r = 0.788, p < 0.01$). Lastly, willingness to sign contracts was highly correlated with perceived value ($r = 0.757, p < 0.01$), see Table 3.6.

Table 3.7. Correlations Study 3.

	1	2	3	4	5
1. Understanding	1				
2. Perc. understanding	0.175^{**}	1			
3. Trust	0.070	0.661^{**}	1		
4. Willingness to sign	0.127^{**}	0.609^{**}	0.838^{**}	1	
5. Perceived value	0.088	0.595^{**}	0.810^{**}	0.780^{**}	1

Note. $N = 414$, ** $p < 0.01$

We ran Pearson product-moment correlations to determine the relationship between contract understanding, perceived contract understanding, trust, willingness to sign contracts and perceived value, see Table 3.7. All correlations were positive. There was a strong correlation between contract understanding and perceived contract understanding ($r = 0.175$, $p < 0.01$), as well as between contract understanding and willingness to sign contracts ($r = 0.127$, $p < 0.01$). There was no correlation between contract understanding and trust in the company, and there was also no correlation between contract understanding and perceived value. Perceived contract understanding was highly correlated with trust in the company ($r = 0.661$, $p < 0.01$), willingness to sign contracts ($r = 0.609$, $p < 0.01$), and value perception ($r = 0.595, p < 0.01$). Furthermore, trust in the company highly correlated with willingness to sign contracts ($r = 0.838$, $p < 0.01$) and perceived value ($r = .810$, $p < 0.01$). Lastly, willingness to sign contracts was highly correlated with perceived value ($r = 0.780, p < 0.01$).

Table 3.8. Correlations Study 4.

	1	2	3	4	5
1. Understanding	1				
2. Perc. understanding	0.065	1			
3. Trust	0.013	0.708^{**}	1		
4. Willingness to sign	-0.044	0.522^{**}	0.634^{**}	1	
5. Perceived value	-0.063	0.540^{**}	0.676^{**}	0.635^{**}	1

Note. $N = 454$, ** $p < 0.01$

We ran Pearson product-moment correlations to determine the relationship between contract understanding, perceived contract understanding, trust in the company, willingness to sign contracts and perceived value. All correlations were positive. There was no correlation between contract understanding and any of the other consumer outcomes. Perceived contract understanding was highly correlated with trust in the company ($r = .708, p < 0.01$), willingness to sign contracts ($r = .552, p < 0.01$) and perceived value ($r = .540, p < 0.01$). Furthermore, trust in the company highly correlated with willingness to sign contracts ($r = .634, p < 0.01$), and perceived value ($r = .676, p < 0.01$). Lastly, willingness to sign contracts highly correlated with perceived value ($r = .635, p < 0.01$), see Table 3.8.

A.2. Study Material

Terms and Conditions

Last updated January 23, 2021

AGREEMENT TO TERMS

These Terms of Use constitute a legally binding agreement made between you, whether personally or on behalf of an entity (“you”) and Coquiptur ("Company", “we”, “us”, or “our”), concerning your access to and use of the Coquiptur website as well as any other media form, media channel, mobile website or mobile application related, linked, or otherwise connected thereto (collectively, the “Site”). You agree that by accessing the Site, you have read, understood, and agreed to be bound by all of these Terms of Use. IF YOU DO NOT AGREE WITH ALL OF THESE TERMS OF USE, THEN YOU ARE EXPRESSLY PROHIBITED FROM USING THE SITE AND YOU MUST DISCONTINUE USE IMMEDIATELY.

Supplemental terms and conditions or documents that may be posted on the Site from time to time are hereby expressly incorporated herein by reference. We reserve the right, in our sole discretion, to make changes or modifications to these Terms of Use at any time and for any reason. We will alert you about any changes by updating the “Last updated” date of these Terms of Use, and you waive any right to receive specific notice of each such change. It is your responsibility to periodically review these Terms of Use to stay informed of updates. You will be subject to, and will be deemed to have been made aware of and to have accepted, the changes in any revised Terms of Use by your continued use of the Site after the date such revised Terms of Use are posted.

Figure 3.9. Manipulation of no information salience—Study 1 and Study 2

Terms and Conditions

Last updated January 23, 2021

AGREEMENT TO TERMS

These Terms of Use constitute a legally binding agreement made between you, whether personally or on behalf of an entity ("you") and Coquiptur ("Company", "we", "us", or "our"), concerning your access to and use of the Coquiptur website as well as any other media form, media channel, mobile website or mobile application related, linked, or otherwise connected thereto (collectively, the "Site"). You agree that by accessing the Site, you have read, understood, and agreed to be bound by all of these Terms of Use. IF YOU DO NOT AGREE WITH ALL OF THESE TERMS OF USE, THEN YOU ARE EXPRESSLY PROHIBITED FROM USING THE SITE AND YOU MUST DISCONTINUE USE IMMEDIATELY.

Supplemental terms and conditions or documents that may be posted on the Site from time to time are hereby expressly incorporated herein by reference. We reserve the right, in our sole discretion, to make changes or modifications to these Terms of Use at any time and for any reason. We will alert you about any changes by updating the "Last updated" date of these Terms of Use, and you waive any right to receive specific notice of each such change. It is your responsibility to periodically review these Terms of Use to stay informed of updates. You will be subject to, and will be deemed to have been made aware of and to have accepted, the changes in any revised Terms of Use by your continued use of the Site after the date such revised Terms of Use are posted.

SUMMARY of the Terms and Conditions

Every company has its terms. These are ours.

We wanted to lay out some important terms and rules before you begin using our platform – we hope you are cool with that!

Here, basically you need to be 18 or above to use the website, and you need to follow the law of your county, state, and the nation that you live in. We will abide by the law, too!

Figure 3.10. Manipulation of information salience—Study 1 and Study 2

Terms and Conditions

Last updated March 15, 2021

AGREEMENT TO TERMS

These Terms of Use constitute a legally binding agreement made between you, whether personally or on behalf of an entity (“you”) and Coquiptur ("Company", “we”, “us”, or “our”), concerning your access to and use of the Coquiptur website as well as any other media form, media channel, mobile website or mobile application related, linked, or otherwise connected thereto (collectively, the “Site”). You agree that by accessing the Site, you have read, understood, and agreed to be bound by all of these Terms of Use. IF YOU DO NOT AGREE WITH ALL OF THESE TERMS OF USE, THEN YOU ARE EXPRESSLY PROHIBITED FROM USING THE SITE AND YOU MUST DISCONTINUE USE IMMEDIATELY.

The Company will use your data—which is anonymized and cannot be traced back to you—in the future.

Supplemental terms and conditions or documents that may be posted on the Site from time to time are hereby expressly incorporated herein by reference. We reserve the right, in our sole discretion, to make changes or modifications to these Terms of Use at any time and for any reason. We will alert you about any changes by updating the “Last updated” date of these Terms of Use, and you waive any right to receive specific notice of each such change. It is your responsibility to periodically review these Terms of Use to stay informed of updates. You will be subject to, and will be deemed to have been made aware of and to have accepted, the changes in any revised Terms of Use by your continued use of the Site after the date such revised Terms of Use are posted.

Figure 3.11. Manipulation of no information salience (here: low contract risk)—Study 3

Terms and Conditions

Last updated March 15, 2021

AGREEMENT TO TERMS

These Terms of Use constitute a legally binding agreement made between you, whether personally or on behalf of an entity ("you") and Coquiptur ("Company", "we", "us", or "our"), concerning your access to and use of the Coquiptur website as well as any other media form, media channel, mobile website or mobile application related, linked, or otherwise connected thereto (collectively, the "Site"). You agree that by accessing the Site, you have read, understood, and agreed to be bound by all of these Terms of Use. IF YOU DO NOT AGREE WITH ALL OF THESE TERMS OF USE, THEN YOU ARE EXPRESSLY PROHIBITED FROM USING THE SITE AND YOU MUST DISCONTINUE USE IMMEDIATELY.

The Company will use your data—which is non-anonymized and can be traced back to you—in the future.

Supplemental terms and conditions or documents that may be posted on the Site from time to time are hereby expressly incorporated herein by reference. We reserve the right, in our sole discretion, to make changes or modifications to these Terms of Use at any time and for any reason. We will alert you about any changes by updating the "Last updated" date of these Terms of Use, and you waive any right to receive specific notice of each such change. It is your responsibility to periodically review these Terms of Use to stay informed of updates. You will be subject to, and will be deemed to have been made aware of and to have accepted, the changes in any revised Terms of Use by your continued use of the Site after the date such revised Terms of Use are posted.

SUMMARY of the Terms and Conditions

Every company has its terms. These are ours.

We wanted to lay out some important terms and rules before you begin using our platform – we hope you are cool with that!

We will use your data—which is non-anonymized and can be traced back to you—in the future.

Here, basically you need to be 18 or above to use the website, and you need to follow the law of your county, state, and the nation that you live in. We will abide by the law, too!

Figure 3.12. Manipulation of information salience (here: high contract risk)—Study 3

Imagine that you are signing up for an account at a social media / cloud storage platform Coquiptur, which allows you to encrypt and secure your personal data using random images.

In compliance with standard legal practice, Coquiptur requires all users to read and agree to its Terms and Conditions. Note that Coquiptur will use your data—which is anonymized and cannot be traced back to you—in the future.

Please click the arrow button below to access the Terms and Conditions. Once you have read through the Terms and Conditions (on the next page), click the blue arrow button to continue.

Figure 3.13. Manipulation of low contract risk—Study 3

Imagine that you are signing up for an account at a social media / cloud storage platform Coquiptur, which allows you to encrypt and secure your personal data using random images.

In compliance with standard legal practice, Coquiptur requires all users to read and agree to its Terms and Conditions. Note that Coquiptur will use your data—which is non-anonymized and can be traced back to you—in the future.

Please click the arrow button below to access the Terms and Conditions. Once you have read through the Terms and Conditions (on the next page), click the blue arrow button to continue.

Figure 3.14. Manipulation of high contract risk—Study 3

LEASE AGREEMENT

THIS LEASE AGREEMENT (hereinafter referred to as the "Agreement") is made and entered into this day of **04/01/2021**, by and between **Park Street Management LLC** (hereinafter referred to as "Lessor") and (hereinafter referred to as "Lessee"). No other tenants are allowed without the written consent of the Lessor, or the execution of a new lease/rental agreement.

A breach of the lease agreement will result in a small ($500) penalty fee.

WHEREAS, Lessor is the landlord of certain real property being, lying and situated in **PA 02134** such real property having a street address of **5130 Parkway, Parkview**. The property is described as follows: **Park Street Property** (hereinafter referred to as "Premises").

WHEREAS, Lessor is desirous of leasing the Premises to Lessee upon the terms and conditions as contained herein; and

WHEREAS, Lessee is desirous of leasing the Premises from Lessor on the terms and conditions as contained herein;

NOW, THEREFORE, the covenants and obligations contained herein and other good and valuable consideration, the receipt and sufficiency of which is hereby acknowledged, the parties hereto hereby agree as follows:

Figure 3.15. Manipulation of no information salience (here: low contract risk)—Study 4

LEASE AGREEMENT

THIS LEASE AGREEMENT (hereinafter referred to as the "Agreement") is made and entered into this day of **04/01/2021**, by and between **Park Street Management LLC** (hereinafter referred to as "Lessor") and (hereinafter referred to as "Lessee"). No other tenants are allowed without the written consent of the Lessor, or the execution of a new lease/rental agreement.

A breach of the lease agreement will result in a large ($5,000) penalty fee.

WHEREAS, Lessor is the landlord of certain real property being, lying and situated in **PA 01234** such real property having a street address of **5130 Parkway, Parkview**. The property is described as follows: **Park Street Property** (hereinafter referred to as "Premises").

WHEREAS, Lessor is desirous of leasing the Premises to Lessee upon the terms and conditions as contained herein; and

WHEREAS, Lessee is desirous of leasing the Premises from Lessor on the terms and conditions as contained herein;

NOW, THEREFORE, the covenants and obligations contained herein and other good and valuable consideration, the receipt and sufficiency of which is hereby acknowledged, the parties hereto hereby agree as follows:

SUMMARY of the Lease Agreement

This lease agreement is between Park Street Management LLC, the landlord ("Lessor"), and you, the tenant ("Lessee").

No other tenant can be part of the agreement without the landlord's written consent.

A breach of the lease agreement will result in a large ($5,000) penalty fee.

The property (premise)'s address is 5130 Parkway, Parkview, PA 01234

The landlord wants to lease the property to you, and you want to lease it from the landlord.

Figure 3.16. Manipulation of information salience (here: high contract risk)—Study 4

In order for a prospective resident to lease the property, the resident needs to read and sign the lease agreement.

A breach of this lease agreement will result in a small ($500) penalty fee.

Please click the arrow button below to access the lease agreement. The blanks in the lease agreement are already pre-filled.

Once you have read through the lease agreement (on the next page), click the arrow button on the next page to continue.

Figure 3.17. Manipulation of low contract risk—Study 4

In order for a prospective resident to lease the property, the resident needs to read and sign the lease agreement.

A breach of this lease agreement will result in a large ($5,000) penalty fee.

Please click the arrow button below to access the lease agreement. The blanks in the lease agreement are already pre-filled.

Once you have read through the lease agreement (on the next page), click the arrow button on the next page to continue.

Figure 3.18. Manipulation of high contract risk—Study 4