ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS
FELLOWS' DISCUSSION PAPER SERIES

BARGAINING WITH ALGORITHMS: AN EXPERIMENT ON ALGORITHMIC PRICE DISCRIMINATION AND CONSUMER AND DATA PROTECTION LAWS

Haggai Porat

Discussion Paper No. 98

06/2024

Harvard Law School
Cambridge, MA 02138

Bargaining with Algorithms: An Experiment on Algorithmic Price Discrimination and Consumer and Data Protection Laws

Haggai Porat[†]

June 2024 Draft

Abstract

Using algorithms to personalize prices is no longer a fringe phenomenon but, rather, the predominant business practice in many online markets, where tracking consumers' every click is the industry standard. Seemingly unrelated, for decades, consumer protection laws have been grounded on the premise that consumers lack meaningful power to bargain over contract terms. This paper suggests that the increasing use of algorithms to set personalized prices based on consumers' behavior opens a path for consumers to "bargain" with algorithms over prices and reclaim market power and for the law to regulate pricing algorithms in light of this interplay between consumers and sellers. To support this, the paper presents the results of a novel pre-registered, incentive-compatible randomized experiment in an online lab setting that tested whether and how consumers bargain with algorithms over price when given the opportunity, by offering participants, in multiple rounds, a \$10 gift card for purchase at a price set by an algorithm based on participants' purchase decisions in preceding rounds. The study further explored the potential for regulating algorithmic pricing with tools from consumer and data protection laws commonly deployed in online consumer markets: a disclosure mandate, the right to prevent data collection ex ante ("cookies laws"), and the right to prevent data retention ex post ("erasure laws" or the "right to be forgotten"). We found clear evidence that participants strategically avoided purchases they would have otherwise made to induce a price decrease in subsequent rounds as they experienced and learned from the algorithmic price changes, albeit not to the rationally-dictated extent. We found that this strategic behavior increased in magnitude and statistical significance in the presence of disclosure, as well as clear evidence that participants offered data protection rights used them strategically: preventing retention or collection of their data in rounds in which they purchased the gift card, so as to prevent a subsequent price increase, and allowing it in rounds in which they declined to purchase, so as to signal a low WTP and benefit from a price decrease in the next round.

[†] Ari J. Zweiman Fellow at the John M. Olin Center for Law, Economics, and Business at Harvard Law School, SJD candidate at Harvard Law School, and PhD candidate at Tel Aviv University School of Economics. For their helpful comments and suggestions, I am grateful to Ronen Avraham, Oren Bar-Gill, Michael Birnhack, Ryan Bubb, Alma Cohen, Kevin Davis, Niva Elkin-Koren, Assaf Hamdani, Scott Hemphill, Marcel Kahan, Kobi Kastiel, Lewis Kornhauser, Tamar Kricheli-Katz, Daphna Lewinsohn-Zamir, Florencia Marotta-Wurgler, Ariel Porat, Kathryn Spier, Stephan Tontrup, Eyal Zamir, Tom Zur, and participants at the 34th Annual Conference of the American Law & Economics Association, the 8th Annual NYU-Penn Empirical Contracts Conference, the 2024 Annual Conference of the Israeli Law & Economics Association, the NYU Law & Economics Workshop, the Hebrew University Law & Economics Workshop, and the Tel Aviv University Law & Tech Workshop. I am also grateful to the John M. Olin Center for Law, Economics, and Business, the Project on the Foundations of Private Law, the Program on Negotiation ("PON Next Generation" grant) at Harvard Law School and the Harvard Institute for Quantitative Social Science (IQSS) for their financial support of this project. The hypotheses of this study were pre-registered in the Penn Wharton Credibility Lab's registry for randomized controlled trials (AsPredicted #157705) on January 11, 2024, and can be found at https://aspredicted.org/blind.php?x=STG_VV8. The author commits to making all data and code publicly available in an online repository following publication, accessible through a link that will be added to this footnote, regardless of whether required by the publishing journal. The experimental design was reviewed and approved (as exempt from full-board review) by the Harvard University IRB on November 9, 2023 (IRB23-1488).

Keywords: Price Discrimination, Algorithmic Pricing, Data Protection, Consumer Protection, Disclosure Mandates, Right to Be Forgotten, Law and Technology, Empirical Legal Studies

JEL Classifications: L11, K12.

1. Introduction

Using algorithms to set the terms of consumer contracts, including the price, is no longer a fringe technological phenomenon but, rather, the predominant business practice in many markets, most notably by platforms like Uber and Amazon. Algorithmic price-setting has various implications, from facilitating price collusion to its impact on consumers' privacy. The implication on which this paper focuses is its enhancement of sellers' ability to set personalized prices, i.e., to price discriminate. Specifically, algorithms are used to estimate consumers' willingness to pay (WTP) for the given product or service with significant precision and in a highly granular fashion. For example, instead of setting one low price for elderly moviegoers and another (higher or "regular") price for everyone else, sellers such as Uber use algorithms to set a potentially infinite number of different prices based on each consumer's WTP, which is estimated based on large volumes of input data fed into the pricing algorithm.

Price discrimination is not a new practice. Before the algorithmic era, it was typically limited to price-setting based on immutable characteristics such as consumer sex, age, or race. However, this focus is fast becoming obsolete, with recent empirical studies showing that consumers' behavior, such as their search, browsing, and purchase histories, is more informative of their WTP than immutable characteristics or group affiliations (see Shiller, 2020). Indeed, personalized pricing based on consumer behavior is becoming increasingly dominant in online markets (see, e.g., Grochowski, Jabłonowska, Lagioia, & Sartor, 2022), where tracking consumers' every click is the industry standard, and minute-to-minute algorithmic price changes are so common that nobody seems to question what exactly drives them.

Seemingly unrelatedly, consumer protection laws have developed for decades based on the premise that consumers lack any meaningful power to bargain over contract terms beyond the ability to switch to a competitor's product. The primary economic rationale has been that high transaction costs lead sellers to efficiently rely predominantly on boilerplate "take it or leave it" contracts. While in some cases prices can be negotiated, such as with small brick-and-mortar businesses, in many other cases prices are mostly set in stone, such as with large chains like

Walmart or Starbucks. This paper suggests that the two features of online markets outlined above – i.e., the increasing use of algorithms to set personalized prices and the reliance on consumers' behavior to draw inferences about their WTP – open a path for consumers to "bargain" with algorithms over prices and exercise market power long deemed to be lost by strategically adjusting their behavior in anticipation of affecting prices. From economic lens, automating the price-setting process has dramatically reduced transaction costs, thereby enabling consumers to bargain with a potentially receptive decision-maker that can easily depart from the take-it-or-leave-it contract and set personalized terms: the algorithm. Contrary to popular sentiments regarding the use of algorithms in consumer markets, this shift happened largely *due* to the algorithm's nonhuman ability to adjust prices at nearly zero transaction costs, not *despite* its nonhuman characteristics.

To be clear, "bargaining" is not meant to imply that consumers can negotiate with algorithms in the simple sense of the term, whereby a consumer uses *words* that signal her low WTP to convince the seller to decrease the price (although that may also be possible soon with the increasing commercial viability of LLMs). Rather, this signaling is achieved by the consumer *behaving* in a way that the algorithm records and interprets as indicative of lower WTP relative to other possible forms of behavior. For example, taking the bus when the Uber fare is too expensive or selecting a cheap brand of toilet paper on Amazon may be regular decisions made by price-sensitive consumers. Alternatively, they could also be the strategic decisions of sophisticated consumers trying to signal a low WTP to the algorithms to affect the prices offered to them (either in the current transaction or in future transactions with the same seller).

Consumers' potential ability to bargain with algorithms over prices entails significant implications for our understanding of how consumer and data protection rights affect consumers. The three legal rights this paper studies are information disclosure mandates, which are widely used in various contexts in both the US and the EU; the right to prevent data collection ex ante ("cookies laws"), which is mandated by the EU's GDPR, but has spilled over to essentially everywhere, including the US (see Marotta-Wurgler & Davis, forthcoming); and the right to prevent data retention ex post ("erasure laws" or the "right to be forgotten"), which is enshrined in both the EU's GDPR (Article 17) as well as in the California Consumer Protection Act (2018). In most jurisdictions, these rights are deployed uniformly across different markets and business practices, applying also to algorithmic price personalization, despite regulators' apparent ignorance to their potential effects on this practice.This paper presents a novel incentivized experimental study to test whether and how consumers bargain with an algorithm over the transaction price given the opportunity, and how different legal rules could regulate this strategic behavior. The experimental design is motivated by a theoretical model of behavior-based price discrimination developed in earlier work (see Porat, 2022). In brief, rational forward-looking consumers who repeatedly engage with the same seller that uses an algorithm to set personalized prices may benefit from strategically avoiding a seemingly beneficial transaction to signal their low WTP (and, hence, trigger a decrease of future prices set by the algorithm) if the lost surplus from the current transaction is more than offset by the increased surplus from future transactions due to the lower prices. Counterintuitively, despite consumers making the best decisions, this strategic behavior may be socially harmful due to efficient transactions not transpiring. Granting consumers legal rights further complexifies the analysis. First, whereas information disclosures are typically meant to help consumers enter transactions only when they are beneficial, disclosing the use of behavior-based algorithmic pricing has a potential effect of intensifying consumers' strategic behavior. Second, Cookies consent law and a right to erasure – rights that accord consumers control over their data – are typically intended to protect consumers' preference for privacy. However, these rights have an overlooked effect – they can be used by consumers to control the information the algorithm can use. Specifically, consumers can strategically decide whether or not to permit data collection to signal their WTP to the algorithm so as to affect prices.

To the end of studying consumers' behavior in the face of algorithmic pricing and its legal regulation, we designed a novel pre-registered, incentive-compatible randomized experiment in an online lab setting where participants, playing a series of up to four rounds, were offered the opportunity to purchase a $10 Walmart gift card for a price set by a pricing algorithm. Specifically, participants were asked in each round they played to decide whether to buy the gift card, with the algorithm setting the price offered in each round based on decisions the participants had made in the preceding round. We supplemented this experimental design with several additional features (presented in section 2.1) to rigorously identify whether consumers learn to avoid purchasing a product not because the price is too high but only in order to signal a lower WTP to the pricing algorithm. The study further explored the potential regulation of algorithmic pricing using three tools from consumer and data protection laws that are commonly deployed in online consumer markets, as detailed above. First, we investigated the effectiveness of regulating consumers' information through disclosure mandates, by explicitly informing some participants that the algorithm would set prices based on their prior purchase decisions. Second, we investigated the effects of data protection tools on consumers' behavior by replicating two prominent legal tools that accord consumers control over their data: the right to prevent data collection ex ante ("cookies laws") and the right to prevent data retention ex post ("erasure laws" or the "right to be forgotten").

Table 1 in section 2.2 provides an overview of all the study's pre-registered hypotheses, as well as whether supporting evidence was found for each hypothesis. In sum, we found clear evidence that some participants strategically avoided purchases they otherwise would have made so as to induce a price decrease by the algorithm in subsequent rounds they played, albeit dramatically less than the rationally dictated extent of strategic avoidance. Specifically, we identify a tug-o-war between experience and strategic incentives: at first, participants learn to behave more strategically as they experience the algorithmic price changes. However, as the remaining number of interactions decreases, participants behave less strategically, as there is less to gain from future price cuts. Using participants' WTP for the gift card, which was elicited in a non-strategic setting, we found some evidence that this behavior came at the cost of avoiding purchases even when the price was lower than the value participants assigned to the gift card. For the treatment groups that were presented with an information disclosure, we found that these effects increased in magnitude and statistical significance. We also found clear evidence that participants who were offered data protection rights used them strategically: deleting their data or preventing its collection in rounds in which they purchased the gift card, so as to prevent a price increase, and allowing their data to be retained in rounds in which they did not purchase the gift card, so as to signal a low WTP to the algorithm and benefit from a price decrease. Lastly, in an exploratory analysis we found some evidence that the ex-post right to delete data was used by participants more effectively than the ex-ante right to prevent the collection of cookies.

To be clear, the finding that consumers can exercise some bargaining power when prices are set algorithmically and that consumer and data protection laws can be used by consumers in ways that go beyond their intended functions is not meant to imply that consumers' strategic behavior is socially desirable from a normative standpoint, nor that regulating algorithmic pricing would yield beneficial outcomes. Indeed, as explained above, there are social costs entailed in consumers' strategic response to algorithmic pricing, captured by the fact that some consumers avoid efficient purchases they would have made if it were not for their effect on future prices. At the same time, this type of algorithmic pricing entails the standard social benefits emanating from price discrimination, i.e., lower prices to consumers with low WTP. A comprehensive theoretical analysis of these (and additional) elements of the welfare tradeoffs has been pursued elsewhere (Porat, 2022). The goal of the current study is purely empirical, i.e., to advance our understanding of the mechanisms through which bargaining with pricing algorithms could occur, and how regulating consumers' information and data may affect this interplay between consumers and algorithms. Hopefully, these findings will inform legal policymaking that is more sensitive than the one currently in place to the specific business practice for which algorithms are used and to the potential effects of consumers and data protection laws, which often extend beyond their intended range.

Prior literature. This paper contributes to several strands of the existing legal and economic literature. First, it complements the theoretical literature on algorithmic price discrimination and behavior-based pricing (e.g., Villas-Boas, 2004; Bar-Gill, 2019; Porat, 2022; Bar-Gill, Sunstein, & Talgam-Cohen, 2023) by empirically testing and validating some of the theoretical assumptions and predictions.. Second, it contributes to the empirical literature on algorithmic price discrimination (e.g., Gillis & Spiess, 2019), and specifically to the scarce experimental studies of this practice (e.g., Hillenbrand & Hippel, 2019; Chen, Bó, & Hakimov, 2023). The paper presents the first attempt to empirically study price discrimination based on consumers' purchasing behavior rather than immutable characteristics, focusing on their strategic responses to the practice. Second, deriving most of its hypotheses from the theory developed in Porat (2022), the study contributes to the literature on the regulation of algorithms, adding to the research of the effectiveness of disclosure mandates in consumer markets in general (e.g., Bakos, Marotta-Wurgler, & Trossen, 2014; Ben-Shahar & Schneider, 2014) and algorithmic transparency in particular (e.g., Kleinberg, Ludwig, Mullainathan, & Sunstein, 2018). Similarly, the paper contributes an important and novel perspective to the literature on data protection rights. Despite its vastness, this body of research focuses almost exclusively on consumer privacy and overlooks the potential of these rights for regulating consumer data and, consequently, business practices that utilize the data, like algorithmic pricing, with only a few studies considering the rights' informational effects (e.g., Marotta-Wurgler (2016); Porat (2024); Taylor (2004); Yoo (2022)).[2]

[2] For a more comprehensive survey of the economic theory of privacy, see Acquisti, Taylor, & Wagman (2016).

The paper proceeds as follows. Section 2 outlines the main features of the experimental design and the pre-registered hypotheses. Section 3 presents the data, various attrition and balance tests we conducted to ensure its validity, and a descriptive exploration of the WTP measure elicited from participants. Section 4 then presents the results regarding the strategic ("bargaining") behavior of the participants in response to algorithmic personalized pricing, while section 5 presents the results regarding the regulation of algorithmic pricing through disclosure, the right to prevent data collection (cookies), and the right to delete data (erasure). Lastly, section 6 wraps up the discussion with some brief concluding remarks. The complete experimental protocol presented to participants is in the Appendix; the pre-registration form can be accessed via a link in footnote 1; and the raw data will be made publicly available following the publication of this paper as supplemental material.

2. Experimental Design and Hypotheses

2.1 Experimental Design

In each of up to four rounds they played, participants were asked whether they wanted to use a portion of a $10 endowment, given to them at the outset of every round, to purchase a $10 Walmart gift card. They were informed in advance of a chance of being awarded a bonus payment (10% of all participants) of what would have resulted from their purchase decision in one of the rounds they played, randomly selected. The participants were informed that an algorithm would set the price of the gift card in each round and that prices might vary across rounds and participants. Specifically, after offering an initial price of $7 to all participants in the first round, the algorithm set the price in consecutive rounds using the following simple recursive rule: an increase of $1 for participants who had made the purchase in the preceding round, a decrease of $1 for those who had declined, and no change in price for participants in two treatment groups given the option to prevent collection of their data and opted to exercise that option. Participants were not explicitly informed about the algorithm's pricing function, but could infer its operation from observing the price changes as the experiment progressed and they gained experience.

To illustrate, if a participant decided not to purchase the $10 gift card in the first round (for $7), their potential bonus payment in that round would be the $10 cash endowment. They would then be offered a gift card for $6 in the second round. If they decided to make the purchase at that price, their potential bonus payment for that round would be the $10 gift card and the remaining

$4 in change (cash). The price in the next round would then go back up to $7, and so on and so forth. Notice that given this algorithm, the price in the fourth round could vary from $4 to $10.

Participants were paid a flat fee of $1.7 to participate in the study, equivalent to $8.3/hour for the median participant (who completed the experiment in 12:16 minutes). In addition, the chance of winning the bonus payment described above, which is a commonly used mechanism for managing the costs of online experiments, almost doubled participants' expected payoff from participating in the study and provided ample incentive to exert effort. Furthermore, as mentioned above, the bonus amount and composition (i.e., $10 in cash or a Walmart gift card and change) were determined by randomly selecting *one* of the rounds the selected participant had played. This meant that participants could receive no more than one Walmart gift card. In experiments involving multiple rounds, this feature is essential for neutralizing a potential "income effect," which alters participants' WTP for the product or service across rounds, particularly in rounds following a purchase. For example, in our study, absent this feature, a participant who had decided to buy the gift card in the first round might have concluded that they now had less of a need for *another* $10 gift card, making it less likely for them to buy the gift card in any of the subsequent rounds even at an identical price.[3] Randomly selecting one of the rounds to determine the final bonus payment neutralized this potential effect by ensuring that participants were incentivized to make the optimal decision in each round as though they had not played any of the previous rounds.

The experimental design incorporated seven experimental conditions, which were divided into two sets of inquiry, with each investigating different aspects of consumers' interaction with algorithmic price discrimination and regulation. The first set of conditions tested consumers' strategic behavior when interacting with algorithmic price discrimination and consisted of four treatment groups that differed by number of rounds participants played (denoted 1R, 2R, 3R, and 4R, with the leading number indicating the number of rounds played). Participants in all the seven treatment groups were informed prior to the first round that the algorithm might offer them different prices in different rounds. However, participants in these four groups were not explicitly

[3] In some scenarios, the opposite may be true, namely, that the second gift card may be more valuable than the first. For example, if Walmart charges shipping fees from customers, it may be beneficial to make more purchases in a single order. Moreover, a second gift card increases the variety of products available for purchase (assuming participants would rather avoid adding additional funds). I thank Florencia Marotta-Wurgler for this insight. For our purposes, this type of income effect is just as problematic and could undermine the validity of our inferences unless neutralized.

told that their purchasing decisions might affect those prices, thereby simulating markets where consumers are aware that prices are set algorithmically but may not realize how their behavior may factor into the price-setting.[4]

The second set of experimental conditions, consisting of three treatment groups, tested the effect of regulating algorithmic price discrimination on consumers' strategic behavior. Participants in these groups played four rounds and received all the information that was conveyed to participants in treatment groups 1R–4R, followed by additional information and instructions, as follows. The first treatment group tested the effect of disclosure on participants' behavior (denoted D) by explicitly informing them that their purchasing decisions in each round could lead to a price change in future rounds. In addition to the impact of this disclosed information, the second and third treatment groups each tested the effect of a different legal rule that affords consumers control over their data. The rule tested by the one group, denoted C, simulated a "cookies regime" that allows consumers to prevent collection of their data: participants were given the option to decide, *prior to* the start of each round, whether to allow the algorithm to use the information about their purchase decision in the current round in future rounds. The rule tested by the second group was the right that their purchase decision be forgotten by the algorithm, simulating an "erasure regime" (denoted E) that allows consumers to request that their data be deleted: participants were presented with an identically phrased option *after* each round to request that their purchase decision be deleted. When a participant in these two groups opted not to allow data collection, the algorithm offered the same price for the gift card in the next round, regardless of the purchase decision. For both treatment groups, participants made their choice when they knew the price in the current round, so that the only difference between the two experimental conditions was the timing of the data-related decision: before or after the purchase decision. In addition, both groups were presented with the choice to prevent collection of data about their decisions only from the second round onward, so that all participants would have experienced a price adjustment at least once before potentially preventing the algorithm from further adjusting prices.

[4] Note that participants may not infer from the instructions that it is necessarily their purchasing decisions that determine the prices set to them, as the instructions remain silent regarding what information is used to set prices, and participants are generally aware that researchers have access to their demographic information, for example. Admittedly, however, informing participants that an algorithm sets the prices does not capture markets where consumers are entirely oblivious to the fact that prices are set algorithmically.

Crucially, in all treatment groups, the participants were informed about the number of rounds they would play only at the start of their *final* round. This crucial design feature allowed us to identify participants' strategic behavior by comparing their purchasing decision in a given round in which they expected to play additional rounds (providing a strategic incentive to take into account the impact of their current decision on future prices) to their decision when they knew they were playing their final round (and thus the decision would not affect future payoffs). To understand the significance of informing participants about the number of rounds they would play only prior to their final round, consider the second-round purchasing decisions made by participants in condition 2R (who knew they were playing their final round) and participants in condition 4R (who expected to play additional rounds). Had the participants known from the outset the number of rounds they would play, the two groups would have diverged in their incentives to purchase the gift card as early as the first round. The reason for this is that awareness of a longer time horizon would have increased participants' strategic incentives to avoid the transaction so as to secure lower prices in subsequent rounds, knowing that their decision in the current round would impact the price in a greater number of future transactions. Accordingly, in this scenario, since groups 2R and 4R would be expected to diverge in their strategic behavior and purchase decisions already in the first round, they would be offered different average prices in the second round, so that comparing their behavior in this round would be like comparing apples to oranges. In other words, the divergence between the groups' first-round incentives would create a selection bias, which would undermine the rigorous identification of participants' strategy and analysis of their decisions in the second, third, and fourth rounds. This risk was neutralized by informing participants about the number of rounds they would play only prior to their final round, in ensuring that all participants reaching any given round had possessed identical information and incentives in all previous rounds.

Figure N2 shows the results of simulations conducted to solve the rational profit-maximizing behavior of consumers. Specifically, for each of the four rounds, the rational decision for each pair of price and the consumer's valuation of the gift card, i.e., the decision the maximizes the sum of expected payoffs from the current and all future rounds, is indicated by the color of the marker: green. if the consumer should purchase the gift card; red, if they should avoid the purchase; or gray, if either decision results in the same expected payoff. The dashed black line indicates the cutoff where price is equal to the value of the gift card. These simulations rely on three assumptions. First, that consumers are risk averse, which is not restrictive given the small amounts at stake. Second, it uses the expected number of rounds remaining conditional on reaching the current round using Bayes rule and the fact that there was an initially equal chance of playing either 1, 2, 3, or 4 rounds, which, admittedly, is not obvious to laypeople. Third, it assumes knowledge of the algorithm's pricing function. Clearly, this is the most restrictive assumption given that participants were *not* told how the pricing algorithm works. However, participants observed the price changes, such that the simulations reflect participants' information better when considering later rounds, and possibly less so for the first round(s). With these caveats in mind, note that in the first round, where the price is $7, *everyone* is better off avoiding the purchase to obtain lower prices in future rounds, even if they value the gift card as high as $10 (i.e., like cash). As one advances through the rounds, it becomes less prudent to avoid a purchase when the price is lower than the value, until the fourth and last round where there is no incentive left to strategically avoid a purchase. Importantly, this exercise demonstrates that the strategic incentive to avoid otherwise efficient purchases decreases over time, which will feature prominently in the results below.

Figure N2: Optimal Purchasing Decisions

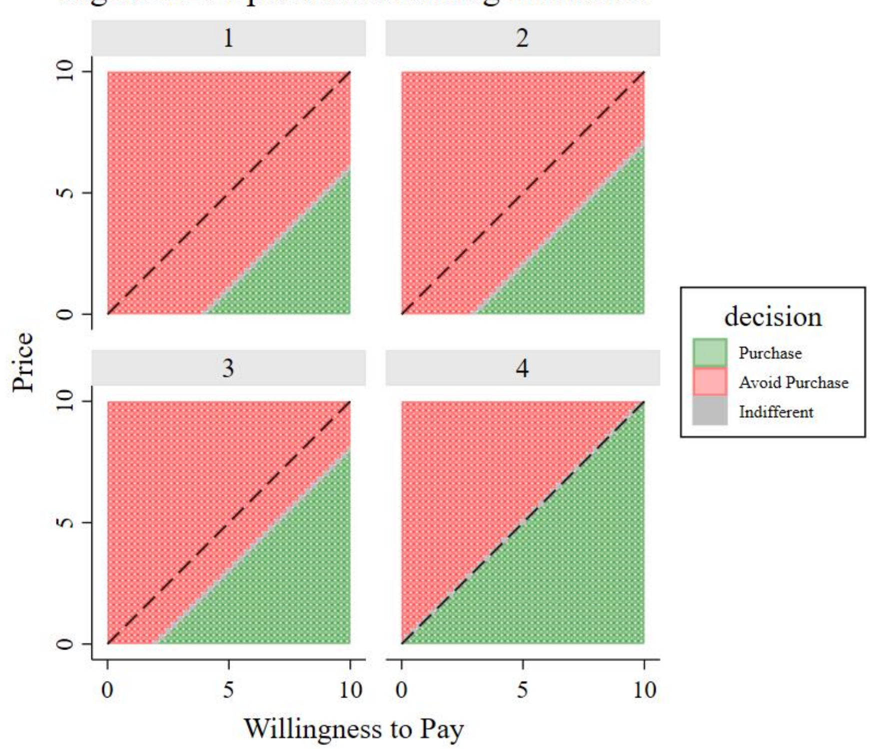

The second and final stage of the experiment was designed to elicit an estimation of all participants' WTP for a $10 Walmart gift card using the Becker-DeGroot-Marschak ("BDM") mechanism (Becker et al., 1964). Playing only a single round in this stage, participants were again given a $10 endowment as a potential bonus and asked to report the highest amount they would be willing to pay for the gift card. Participants were informed that the price of the gift card would not be determined based on their reported WTP, which would have led rational participants to report a lower amount than they were actually willing to pay, but, rather, by a random computer-drawn integer between 0 and 10. If the drawn number was higher than a participant's reported WTP, the gift card would not be purchased and her potential bonus would be $10 in cash. If the drawn number was equal to or lower than the reported WTP, the gift card would be purchased for a price equal to that number. Thus, from a mechanism-design perspective, the BDM mechanism is expected to induce truthful reporting, since the reported WTP affects whether a purchase will occur but not the price.[5]

In addition, participants were required to complete up to three comprehension tests at several stages of the experiment to verify that they understood the instructions and mechanisms of the different procedures. These included one "main" comprehension test of the experiment's general instructions, administered to all participants after having been presented with the instructions, which we used to exclude participants who failed the test from the analysis (as detailed in Section 3); a second comprehension test presented to participants in treatment groups D, C, and E, which assessed participants' understanding of the information disclosure[6]; and a third comprehension test, administered to all participants during the second stage of the experiment, which participants could not fail but were instead forced to answer correctly to be able to proceed to the following screen, which ensured that participants would read the instructions regarding the

---

[5] To see why a participant can do no better than report their true WTP, consider a hypothetical participant whose true WTP for the gift card is $6. First, in a scenario where the drawn number is 7 or higher, the participant can do no better than report her true WTP and prevent a purchase, because reporting a higher figure might result in paying more than the gift card's worth to her, and underreporting will yield the same outcome as reporting the true WTP, where no purchase occurs. Second, in a scenario where the drawn number is 6 or lower, the participant can do no better than report her true WTP and purchase the gift card for a price equal to the drawn number, because reporting a lower WTP might result in missing out on a beneficial transaction without any offsetting benefit since the reported WTP has no effect on the price. At the same time, reporting a higher WTP will yield the same outcome as reporting the true WTP, where the purchase will occur at the same price. Thus, since participants can do no better than report their true WTP regardless of the number drawn, this mechanism induces truth-telling.

[6] Which was *not* used for exclusion of participants from the analysis, as this test was administered to only some of the participants, so that using it to exclude participants would risk introducing a sample selection bias. Rather, the test was administered to ascertain and validate the comprehension level of the information disclosure.

BDM elicitation procedure of their WTP. In a post-experimental survey, participants were asked to provide basic demographic information, to respond to several questions about their online shopping habits, and to describe their perceptions of the study's purpose. In addition, data was collected on participants' geographic location, IP address, and the amount of time they spent on each screen. The experimental protocol that was shown to participants is in the Appendix.

2.2 Hypotheses

Algorithms are increasingly being used to set prices in consumer markets, with sellers using highly granular consumer data to price-discriminate more accurately. However, as discussed in the Introduction, the primary source of this highly granular data is consumers' behavior, such as browsing activity, "likes," and purchase history, rather than immutable characteristics such as sex and age. This shift has created an often overlooked opportunity for consumers to strategically adapt their behavior to affect the prices they are offered for services and products, which is the functional equivalent to bargaining with the algorithm over the contract terms. The theoretical literature surveyed above generates several hypotheses about consumers' behavior and the effects of regulating algorithmic pricing that our study set out to test.

The eight hypotheses presented in Table 1 below (along with other design features of this study) were derived primarily from the theory developed in Porat (2022). They were pre-registered in the Penn Wharton Credibility Lab's registry for randomized controlled trials and are publicly available.[7] Hypotheses 1 and 2 relate to participants' strategic behavior when interacting with algorithmic price discrimination; hypotheses 3 and 4 relate to the potential use of disclosure mandates for regulating algorithmic price discrimination; and hypotheses 5 to 8 relate to the potential use of cookies and erasure laws to regulate algorithmic price discrimination.

Note that the table also previews the study's findings, stating in brief whether evidence was found in support of each hypothesis. Of course, the statements regarding supporting evidence are not objective facts, nor could they be formulated in a binary term, especially given the inclusion of multiple rounds in the experiment. For example, we consistently found no evidence of any of the hypothesized outcomes in the first round of the experiment, which demonstrates that

[7] AsPredicted #157705, registered on January 11, 2024 (https://aspredicted.org/blind.php?x=STG_VV8). To be clear, the wording in Table 1 is not identical to the formulations in the pre-registration form. However, any discrepancies were motivated solely by expositional and proofing considerations. Readers are encouraged to examine the pre-registration form to ascertain that the hypotheses are reported accurately and truthfully here in all meaningful senses.

participants needed to interact with and experience the price-adjustment function of the algorithm at least once to learn how to behave strategically, as opposed to figuring this out solely by reading the experiment's instructions and engaging in game-theoretic reasoning. This learning-from-experience effect was not a pre-registered hypothesis, as we could not theoretically rule out the possibility that participants would engage in strategic behavior to *some* degree already in the first round. However, we were not surprised by the finding and, indeed, anticipated the effect of experience in designing certain features of the experiment. For example, this anticipated effect underlay the choice to allow participants to prevent data collection only from the second round. In sum, the preview of the findings in Table 1 is meant to provide a crude overview of the *arguments* that the statistical analysis in Sections 4 and 5, below, seeks to validate.

Table 1: Hypotheses

Topic	Hypothesis	Evidence
Bargaining with Algorithms (conditions 1R–4R)	H1: Consumers avoid purchases in anticipation of future prices	clear
	H2: Consumers avoid efficient transactions in anticipation of future prices	some
Regulating Consumer Information (condition Disclosure)	H3: Disclosure increase consumers' strategic avoidance of purchases	clear
	H4: Disclosure increases strategic avoidance of efficient purchases	clear
Regulating Data Protection (conditions Cookies & Erasure)	H5: Consumers make different use of data protection rights when invoked before (cookies) versus after (erasure) a transaction	some
	H6: Consumers are more likely to allow data collection/retention when not purchasing	clear
	H7: Data protection rights decrease strategic avoidance of purchases	some
	H8: Data protection rights decrease strategic avoidance of efficient purchases	weak

Note: evidence strength is reported on the following non-cardinal scale: none, weak, some, and clear.

3. Data

This Section will first present the data sampling approach, results of an attrition and balance tests, the (pre-registered) exclusion criteria, and a description of the resulting sample. It will then move on to present the data on participants' elicited WTP and several validation tests.

We implemented the experimental design outlined above using Qualtrics and recruited participants using Prolific. Participation was open to adult (aged 18 and up) US citizens who were fluent in English. Sampling was further restricted to eligible individuals in the highest bracket of approval rate on Prolific (namely, their submissions are rarely rejected by researchers, to avoid potential bots and "spammers") and who do not participate in Prolific with a partner or a friend (to avoid potential cross-treatment contamination). As noted, participants were paid a flat fee of $1.7 to participate in the study, equivalent to $8.3/hour for the median participant (who completed the study in 12:16 minutes), in addition to the potential of a bonus payment in the form of additional money and a Walmart gift card, awarded to a random 10% of participants who completed the experiment and passed the main comprehension test.

To reach the pre-registered sample size of ~200 participants for each of the seven experimental conditions who passed the comprehension test and completed the experiment, we recruited 1,880 individuals who met the above-noted requirements. Of these, 281 (14.95%) did not complete the experiment, and 7 (0.37%) were identified by Qualtrics as bots, resulting in 1,592 valid participants who submitted complete responses (for now, including those who failed the main comprehension test). Table 2, below, presents the results of three attrition tests used to ascertain that attrition rates were similar across experimental conditions, thereby implying that attrition was not caused by features that differed across the conditions. Column 1 reports the results of a logistic regression where the dependent variable indicates whether the experiment was completed and the seven independent variables indicate the assigned condition. Surprisingly, we found that participants assigned to the Cookies condition were somewhat more likely to complete the experiment ($p < 0.01$). To test whether this was caused by something in the design of the Cookies condition that potentially biased the sample or whether it was an inconsequential statistical anomaly, we ran a similar logistic regression (whose results are reported in column 2) where, now, the dependent variable indicates whether a participant reached the stage where the experiment diverged for different treatment groups.[8] The results indicate a lower attrition rate for the Cookies condition even before the experiment diverged across conditions in any way observable to participants, meaning that attrition up to that point did not introduce any sampling bias. Moreover, column 3 reports the results of an attrition test, similar to the first test but including only participants who have reached the point where the conditions diverged, indicating that attrition rates were statistically indistinguishable across all conditions. The results of this decomposition are consistent with the fact that more than half (148) of the participants who did not complete the experiment exited before completing the main comprehension test and observing any condition-specific content.

Table 2: Logistic Regressions Predicting Completion of the Experiment by Assigned Condition

	(1) Completed	(2) Condition divergence	(3) Completed conditional on divergence
1R	0	0	0
	(.)	(.)	(.)
2R	.029	.043	-.012
	(.03)	(.024)	(.023)
3R	.0064	.045	-.038
	(.032)	(.024)	(.025)
4R	-.022	.0085	-.033
	(.032)	(.026)	(.024)
Disclosure	-.00059	.032	-.033
	(.032)	(.025)	(.024)
Cookies	$.079^{**}$	$.056^{*}$	.028
	(.029)	(.024)	(.02)
Erasure	.019	$.051^{*}$	-.03
	(.031)	(.023)	(.024)
Observations	1876	1876	1732
Mean of dep. Var.	.85	.92	.92

Standard errors in parentheses.
Table reports marginal effects.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The study's final sample comprised 1,207 participants, resulting from the exclusion of 385 (24.18%) of the 1,592 valid participants who failed the main comprehension test, as prescribed by

[8] This divergence began after participants completed the main comprehension test, at which stage participants in the Disclosure, Cookies and Erasure conditions were presented with the information disclosure, whereas the participants assigned to the other four conditions were not.

the pre-registered exclusion criteria. The test questions, which are in the Appendix, checked participants' understanding of the pricing mechanism, the potential bonus payment, and the features of the Walmart gift card. Because the rate of participants who failed the test was somewhat higher than expected, the final sample was slightly smaller than planned, with 170 to 176 participants in each treatment group. At the same time, however, the effective screening achieved by the comprehension test bolstered the reliability of the participants who passed it, who, on average, spent 2:18 minutes reading the experiment's instructions and an additional 1:15 minutes completing the comprehension test. Table 3 below presents some descriptive statistics about the behavior and characteristics of the participants comprising the final sample, differentiated by assigned experimental condition. While this sample is not meant to be a representative sample of the US population, it is noteworthy that the median age of Americans is 39 whereas the median age of participants is 38, and that 71% of Americans are white compared to 67% of the participants. In contrast, the are somewhat more female than male participants (58%), and participants are more educated and have a lower income than the average American. However, Amazon users, for example, are also more educated and more female, and this may be a more relevant population in the context of studying algorithmic pricing in online markets.

Table 3: Descriptive Statistics

Condition:	All		1-Round		2-Rounds		3-Rounds	
	Mean (sd)	Min/Max	Mean (sd)	Min/Max	Mean (sd)	Min/Max	Mean (sd)	Min/Max
Round 1 Purchase	.3 (.46)	0 / 1	.34 (.48)	0 / 1	.32 (.47)	0 / 1	.34 (.47)	0 / 1
Round 1 Price	7 (0)	7 / 7	7 (0)	7 / 7	7 (0)	7 / 7	7 (0)	7 / 7
Round 2 Purchase	.3 (.46)	0 / 1			.45 (.5)	0 / 1	.31 (.46)	0 / 1
Round 2 Price	6.6 (.91)	6 / 8			6.6 (.93)	6 / 8	6.7 (.95)	6 / 8
Round 3 Purchase	.45 (.5)	0 / 1					.52 (.5)	0 / 1
Round 3 Price	6.1 (1.6)	5 / 9					6.3 (1.7)	5 / 9
Round 4 Purchase	.49 (.5)	0 / 1						
Round 4 Price	5.9 (1.9)	4 / 10						
Willingness to Pay	5.1 (2.7)	0 / 10	4.8 (2.6)	0 / 10	5 (2.8)	0 / 10	5.2 (2.7)	0 / 10
Response time (minutes)	14 (6.6)	3 / 59	13 (6.8)	5.2 / 38	14 (6.4)	4.8 / 36	13 (6.7)	3.8 / 42
Female	.58 (.49)	0 / 1	.56 (.5)	0 / 1	.57 (.5)	0 / 1	.57 (.5)	0 / 1
Age	41 (14)	18 / 88	40 (14)	20 / 77	42 (15)	19 / 81	40 (13)	19 / 75
(Non-Hispanic) White	.67 (.47)	0 / 1	.67 (.47)	0 / 1	.71 (.45)	0 / 1	.67 (.47)	0 / 1
Income above $50k/year	.46 (.5)	0 / 1	.46 (.5)	0 / 1	.45 (.5)	0 / 1	.49 (.5)	0 / 1
College education or higher	.58 (.49)	0 / 1	.6 (.49)	0 / 1	.56 (.5)	0 / 1	.59 (.49)	0 / 1
Observations	1207		170		173		171	

Condition:	4-Rounds		Disclosure		Cookies		Erasure	
	Mean (sd)	Min/Max	Mean (sd)	Min/Max	Mean (sd)	Min/Max	Mean (sd)	Min/Max
Round 1 Purchase	.27 (.45)	0 / 1	.32 (.47)	0 / 1	.25 (.43)	0 / 1	.26 (.44)	0 / 1
Round 1 Price	7 (0)	7 / 7	7 (0)	7 / 7	7 (0)	7 / 7	7 (0)	7 / 7
Round 2 Purchase	.27 (.44)	0 / 1	.25 (.43)	0 / 1	.25 (.43)	0 / 1	.26 (.44)	0 / 1
Round 2 Price	6.5 (.89)	6 / 8	6.6 (.93)	6 / 8	6.5 (.87)	6 / 8	6.5 (.88)	6 / 8
Round 3 Purchase	.45 (.5)	0 / 1	.41 (.49)	0 / 1	.37 (.48)	0 / 1	.49 (.5)	0 / 1
Round 3 Price	6.1 (1.6)	5 / 9	6.1 (1.6)	5 / 9	6 (1.4)	5 / 9	6 (1.4)	5 / 9
Round 4 Purchase	.43 (.5)	0 / 1	.5 (.5)	0 / 1	.53 (.5)	0 / 1	.51 (.5)	0 / 1
Round 4 Price	6 (2.1)	4 / 10	5.9 (2)	4 / 10	5.7 (1.8)	4 / 10	5.9 (1.8)	4 / 10
Willingness to Pay	5.3 (2.6)	0 / 10	5.4 (2.6)	0 / 10	5 (2.6)	0 / 10	5 (2.7)	0 / 10
Response time (minutes)	14 (7)	4.8 / 49	14 (7)	4.4 / 59	14 (5.9)	4.3 / 38	14 (6.1)	3 / 35
Female	.65 (.48)	0 / 1	.56 (.5)	0 / 1	.56 (.5)	0 / 1	.61 (.49)	0 / 1
Age	42 (15)	18 / 88	40 (14)	19 / 75	39 (13)	18 / 76	42 (15)	18 / 79
(Non-Hispanic) White	.7 (.46)	0 / 1	.63 (.48)	0 / 1	.68 (.47)	0 / 1	.65 (.48)	0 / 1
Income above $50k/year	.47 (.5)	0 / 1	.49 (.5)	0 / 1	.39 (.49)	0 / 1	.45 (.5)	0 / 1
College education or higher	.61 (.49)	0 / 1	.59 (.49)	0 / 1	.53 (.5)	0 / 1	.57 (.5)	0 / 1
Observations	172		171		176		174	

Lastly, we conducted a series of balance tests to ensure that the randomization resulted in similar groups as measured by all the observed characteristics. The results are presented in Table 4, where each column reports the results of a regression of one characteristic on the seven dummy variables for each of the assigned experimental conditions. The results reported in columns 1–5 show that the groups were statistically indistinguishable in terms of participants' sex, age, race, income, and education. Column 6 reports results showing that participants assigned to the Erasure and Cookies conditions took slightly longer than other groups to complete the experiment, which is not surprising given that they had an additional decision to make in each round. Lastly, column 7 presents the results of our test of divergences in participants' WTP for the gift card, which was elicited in the final stage of the experiment using the BDM mechanism described in subsection 2.1. We found that the average WTP in the Disclosure condition ($5.4) was slightly higher than the overall average WTP ($5.1) ($p = 0.03$).

Table 4: Balance Tests

	(1) Female	(2) Age	(3) White	(4) Income above $50k/year	(5) College degree	(6) Time to complete (mins)	(7) WTP
1R	0	0	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
2R	.0018	1.7	.04	-.02	-.039	.48	.24
	(.054)	(1.5)	(.05)	(.054)	(.053)	(.71)	(.29)
3R	.0084	-.4	-.0039	.021	-.0094	-.29	.41
	(.054)	(1.5)	(.051)	(.054)	(.053)	(.71)	(.29)
4R	.086	1.5	.027	.0062	.01	1.1	.5
	(.053)	(1.5)	(.05)	(.054)	(.053)	(.71)	(.29)
Disclosure	-.0092	.36	-.039	.027	-.0094	1.1	.63*
	(.054)	(1.5)	(.052)	(.054)	(.053)	(.71)	(.29)
Cookies	-.0022	-1.3	.0055	-.078	-.072	1.3	.24
	(.053)	(1.5)	(.05)	(.053)	(.053)	(.71)	(.29)
Erasure	.05	1.8	-.021	-.011	-.025	1.4*	.24
	(.053)	(1.5)	(.051)	(.054)	(.053)	(.71)	(.29)
Observations	1207	1207	1207	1207	1207	1207	1207
Mean of dep. var.	.58	41	.67	.46	.58	14	5.1

Standard errors in parentheses.
Marginal effects from logistic regression results reported in columns 1, 3-5.
OLS regression results reported in columns 2, 6, & 7.
Constant estimated in all models but not reported.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Identifying participants' WTP involved both practical and conceptual challenges. The conceptual challenge relates to whether the BDM mechanism can be relied on to accurately measure individuals' WTP as is theoretically expected. The most obvious issue is that the mechanism assumes that individuals can understand the game-theoretic rationale (described in footnote 5) that renders truthful reporting of one's WTP beneficial. To contend with this potential drawback, we included a comprehension test (see the Appendix) prior to eliciting participants' WTP to ensure participants' understanding of the BDM mechanism, comprised of four questions about the procedure's outcome in four hypothetical scenarios that participants had to answer correctly to proceed with the experiment.[9]

The practical challenge is that it is uncertain that individuals know their own WTP even if they want to report it truthfully. In the concrete circumstances of our study, accurate identification of participants' WTP might be undermined by the fact that different experimental conditions might have impacted participants' perception of their WTP in different ways. For example, participants assigned to conditions with more rounds spent more time and cognitive resources determining whether they should purchase the gift card for the price offered each round.

To contend with the second challenge, we plotted the distributions of participants' reported WTP for each experimental condition, presented in Figure 1 below, with an added plot—for eyeball comparison—of a normal distribution taking the mean and variance of participants' WTP in the entire sample. A spike is apparent in the bin where the WTP is equal to 5, **i.e.,** the middle option between 0 and 10, indicating that participants found it challenging to figure out their own WTP and consequently resorted to heuristics. Even if not perfectly precise, this is a sufficiently effective measure of the WTP for our purposes, as demonstrated by the four bottom histograms

[9] Since this test was administered in the final stages of the experiment, it was intended only to induce comprehension of the BDM mechanism rather than exclude participants who failed to understand it. Accordingly, participants were not allowed to begin the procedure before correctly answering all four questions. As expected, participants initially spent 1:57 minutes on average reading the procedure instructions and then 2:23 minutes successfully responding to all four comprehension questions, with 45.4% of participants taking advantage of the option to return to the previous screen to re-read the instructions.

that show that participants reported a higher WTP when they purchased the gift card for $7 in the first round or if in the post-experimental questionnaire, they responded that they shop at Walmart at least once a month on average. Furthermore, despite the finding that emerges from column 7 in Table 4, above, that participants in the disclosure condition reported a slightly higher *average* WTP than other conditions, we found no statistically significant divergences between the conditions in reported WTP when testing between entire *distributions*, thus bolstering our confidence in this measure. Table 5, following Figure 1, reports the p-values obtained from 21 Kolmogorov–Smirnov tests of equality of the WTP distributions of all possible pairwise combinations of experimental conditions.

Figure 1: Willingness to Pay

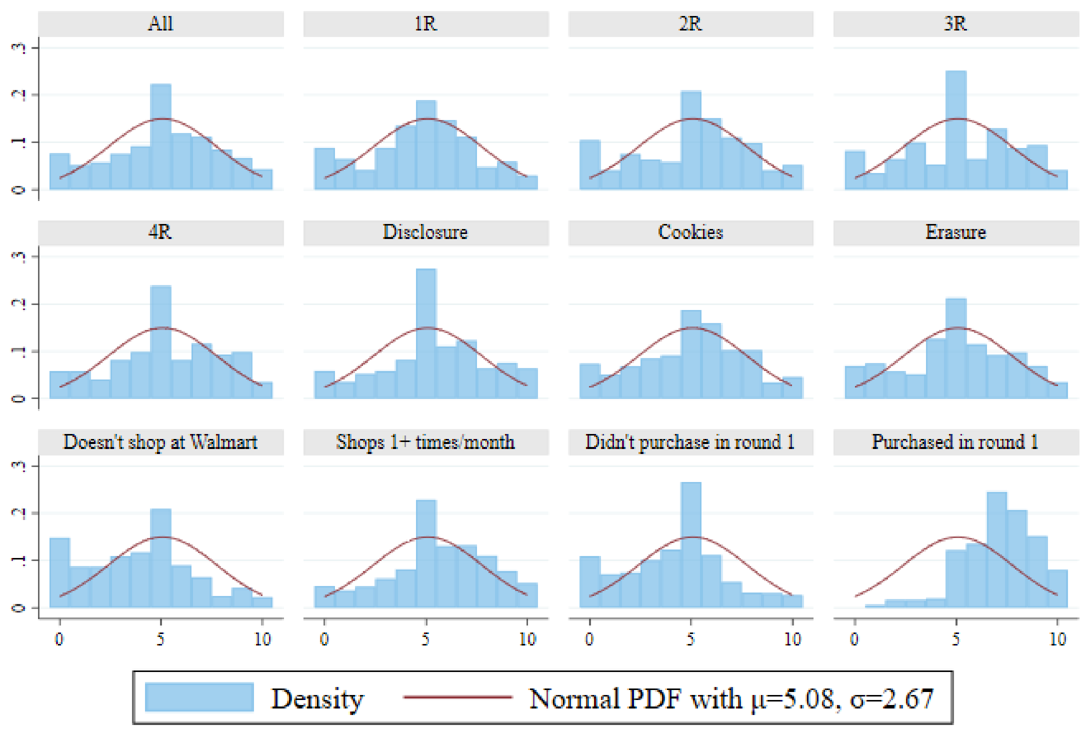

Table 5: Pairwise Tests for Equality of WTP Distributions

Condition	1R	2R	3R	4R	Disclosure	Cookies	Erasure
1R	1.00						
2R	0.66	1.00					
3R	0.28	0.97	1.00				
4R	0.38	0.85	1.00	1.00			
Disclosure	0.10	0.62	0.71	0.97	1.00		
Cookies	0.98	1.00	0.80	0.90	0.55	1.00	
Erasure	0.82	0.99	0.91	0.97	0.41	1.00	1.00

p-values from pairwise Kolmogorov–Smirnov tests of the WTP distributions

4. Bargaining with a Pricing Algorithm

As noted in sub-section 2.2, we found clear evidence that participants strategically adjusted their behavior by avoiding purchasing the Walmart gift card, so as to signal a lower than their actual WTP to the algorithm and be offered lower prices in subsequent rounds. Before presenting this finding, a descriptive presentation of the data can be instructive. Figure 2 below presents a bubble graph where each circle represents the relative number of participants who either purchased the gift card (the red circles) or did not purchase the gift card (the blue circles) for each pair of possible prices and WTP. The black dashed line delineates the points where the prices offered were equal to the WTP, such that efficient transactions are located above the dashed line and inefficient transactions below it.[10] If the WTP was accurately reported and participants did not engage in strategic behavior, we would expect all circles above the line to be red and all circles below the line to be blue. However, while there is a visible partition between blue and red circles, it is certainly not clear-cut. Specifically, the blue circles *above* the dashed line represent participants who may have strategically refrained from an efficient transaction to affect future prices. However, these could also be participants who made an error in their decision not to purchase or whose reported WTP is not accurate, consistent with the well-established theory that consumers' exhibited preferences are affected by the context in which their decisions are made (see Bettman, Luce, & Payne, 1998). Indeed, the red circles below the dashed line, which represent participants who purchased the gift card despite valuing it below its price, cannot be rationally explained.

[10] We use the term "efficient" to refer to the socially first-best outcome, where consumers purchase a product if and only if they value it above its price. This abstracts away from the fact the costs of production may be lower than the price. More importantly, it is not meant to imply that consumers *should* purchase the gift card if it is efficient to do so. The main theme in this study is that it may be "efficient" for consumers not to purchase the gift card even if the price is lower than their WTP. Since using the term efficient to describe both situations risks confusion, we refer to those participants who strategically (and possibly prudently) avoiding an efficient transaction as behaving inefficiently.

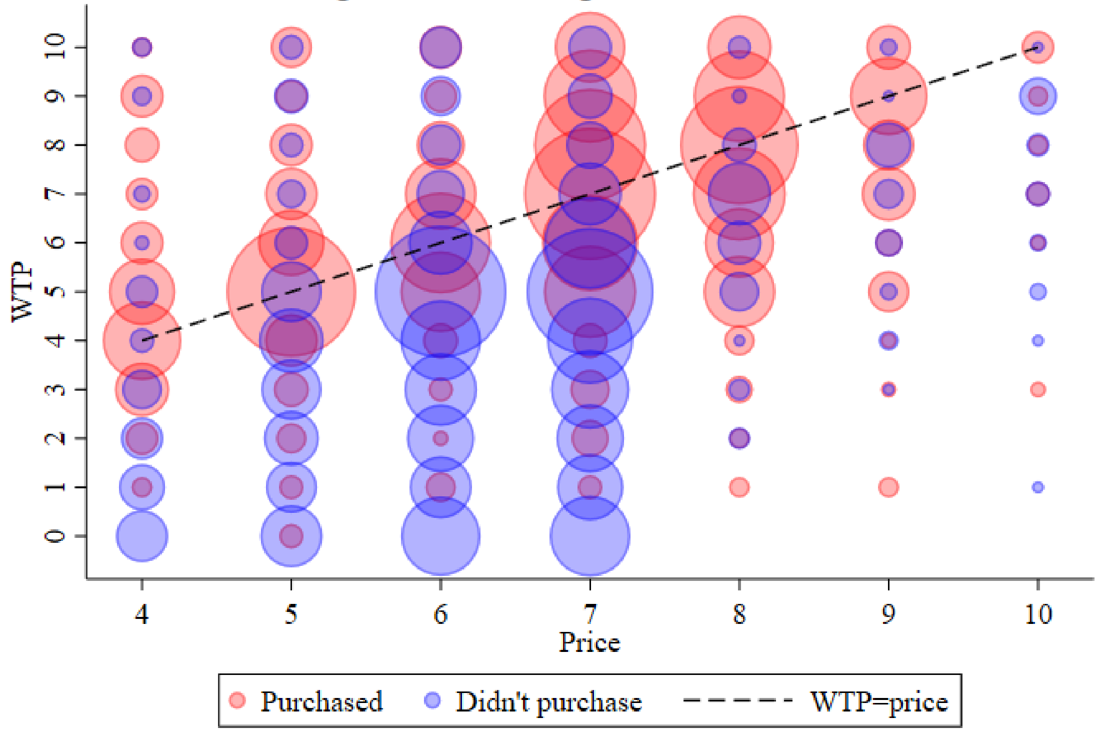

Figure 2: Purchasing decisions and WTP

Figure 3 further reinforces the cautious claim that we are effectively (albeit not perfectly) identifying participants' true WTP by plotting the percentage of participants who purchased the gift card as a function of the transaction's surplus (i.e., the difference between the WTP and the price), where the size of the circle indicates the relative number of participants in that category (which allows us to sidestep outliers such as the single participant who purchased the gift card despite valuing it $8 below its price). Most importantly, an apparent spike is observed when the transactional surplus becomes nonnegative, implying that there is no significant noise in participants' reports of their WTP when measured against their actual behavior. Note that while it is tempting to draw further inferences from the slopes of the graph, this would be imprudent as they are biased downwards by construction.[11]

[11] To see this, note that the transactional surplus is negatively correlated with price, meaning that the observations with higher surplus tend to come from transactions with lower prices. Furthermore, recall that the pricing algorithm reduces the price of the gift card following a participant's decision not to purchase it. Taken together, this means that transactions with higher surplus over-represent participants who had previously refrained from purchasing the gift card, making them less likely to purchase at the lower price compared to the average participant, which biases the slope of the graph downwards. In other words, while surplus affects the purchasing decision, the purchasing decision, in turn, affects the surplus in future rounds.

Figure 3: Purchasing decisions and surplus

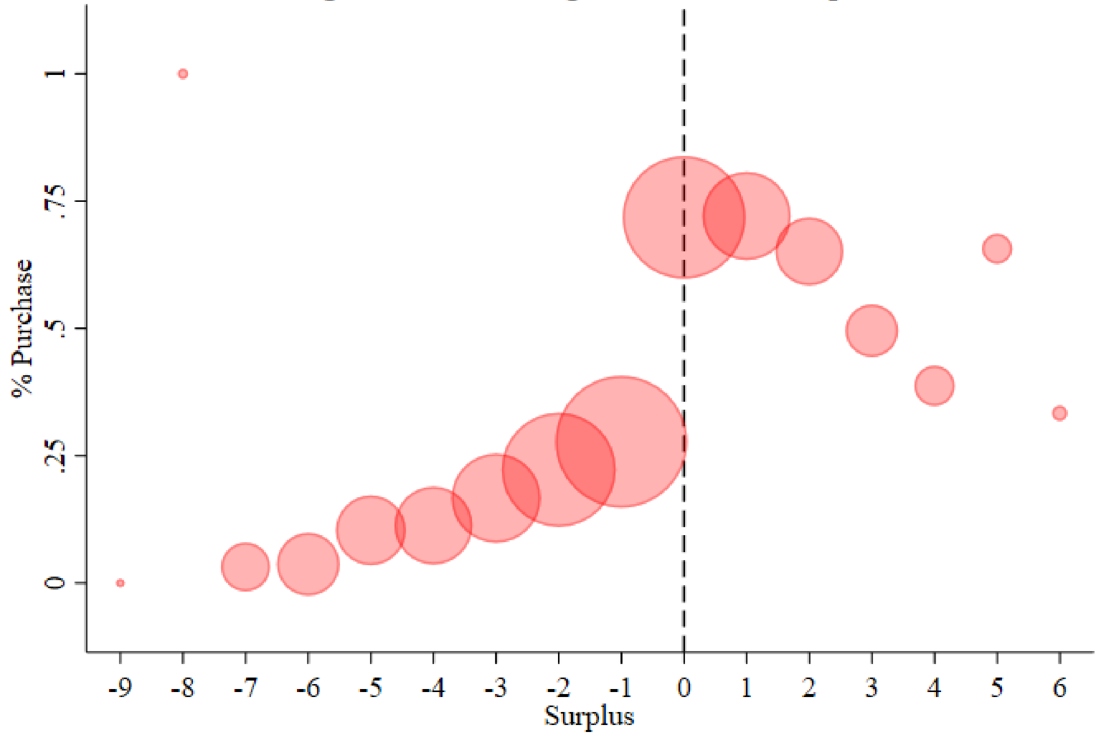

To test Hypothesis 1 (H1) and identify whether participants strategically avoided purchasing the gift card to signal their WTP to the pricing algorithm, we compared, focusing on conditions 1R, 2R, 3R, and 4R, the purchasing decisions of participants for whom any given round was their final round with the decisions of participants who expected to play at least one more round. The identification assumption is that participants playing the final round would have no reason to consider anything other than the price offered to them and their WTP, whereas participants expecting to play additional rounds might have benefitted from strategically refraining from purchasing the gift card so as to secure lower future prices. Note that the fact that these participants were informed about the total number of rounds they would play only prior to their final round ensured that the comparison of purchase decision in each round is between groups of participants who had possessed identical information and, hence, made the same decisions (on average) in all rounds preceding the round being analyzed.[12]

[12] Indeed, we confirmed that the purchasing decisions of participants assigned to conditions 2R, 3R, or 4R in round 1, and that of participants assigned to conditions 3R or 4R in round 2, were statistically indistinguishable ($p = 0.4$ and $p = 0.39$, respectively). Note that all reported p-values in the paper, unless otherwise indicated, are derived from a chi-square test and are robust to a Fisher's exact test.

Figure 4, below, plots, for each round, the percentage of participants who purchased the gift card when it was their final round (in green), the percentage of participants who made the purchase when it was not their final round (in blue), and the difference between these two groups (in red). As Figure 4 demonstrates, we found no divergences in purchasing decisions in the first round, indicating that participants were unable to conclude from the experiment's instructions alone that they may be better off not purchasing the gift card for $7 in the first round. However, in the second round, once participants had engaged with the pricing algorithm and observed the price change (either upward or downward), we found that expecting to play an additional round significantly reduced the likelihood of their purchasing the gift card ($p < 0.01$), indicating that some participants were strategically signaling their WTP to the pricing algorithm. We did not have sufficient statistical power to test effect in the third round. Indeed, while we observed in the third round that participants expecting to play an additional round were less likely to purchase the gift card than participants playing their final round, this divergence was statistically insignificant ($p = 0.21$). This can be explained by the fact that participants assigned to condition 4R knew that only one round remained after the third round, which significantly decreased the strategic incentives to avoid a purchase, as only a few participants of those in condition 4R could potentially gain from this behavior: specifically, whereas all participants would benefit from avoiding a purchase in the first round, only 60 participants in condition 4R in round 3 could benefit from strategically avoiding a purchase in anticipation of a fourth round.[13]

[13] Since avoiding a purchase results in forfeiting the surplus in the current round in exchange for increasing the potential surplus by $2 in the following round, the only scenario in which it is beneficial for a participant to avoid a purchase in the third round given the prospect of playing one additional round is if the price offered renders her surplus from the transaction equal to either $0 or $1. As noted, there were only 60 such participants in condition 4R in round 3.

Figure 4: Strategic Purchasing

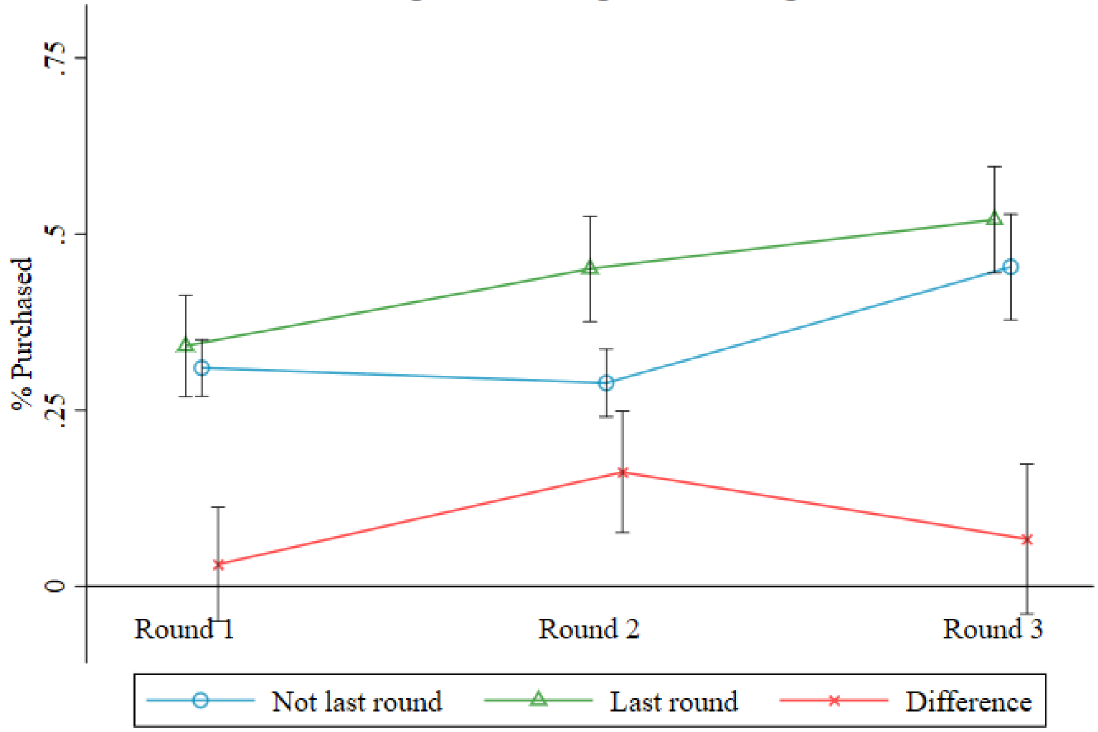

While the experimental design ensured that all participants in the 1R-4R treatment groups possessed the same information and purchase incentives up to the last round they played, it might be argued that the results presented in Figure 4 could have been affected by the slightly divergent (albeit statistically insignificant) first-round purchase decisions made by these participants and, consequently, the different second-round prices they were offered. As a robustness check, Figure 5 plots participants' purchasing decisions for every possible branch of purchase (hence price) history across the three rounds. Specifically, a marker following a solid line represents participants who purchased the gift card in the preceding round, while a marker following a dashed line represents participants who did not purchase the gift card in the preceding round. Note that round 3 is plotted for transparency's sake, but there was virtually no statistical power to analyze participants' behavior in that round, as most of these branches represent groups as small as 10, 13, and 14 participants. Focusing on the second round, we observed that the results found in this round are, indeed, robust to separating participants based on their purchase history. Interestingly, the results are driven more by participants who did not purchase the gift card in the first round ($p <$

0.01), which is consistent with the hypothesis that strategically minded consumers are more likely to avoid purchases.

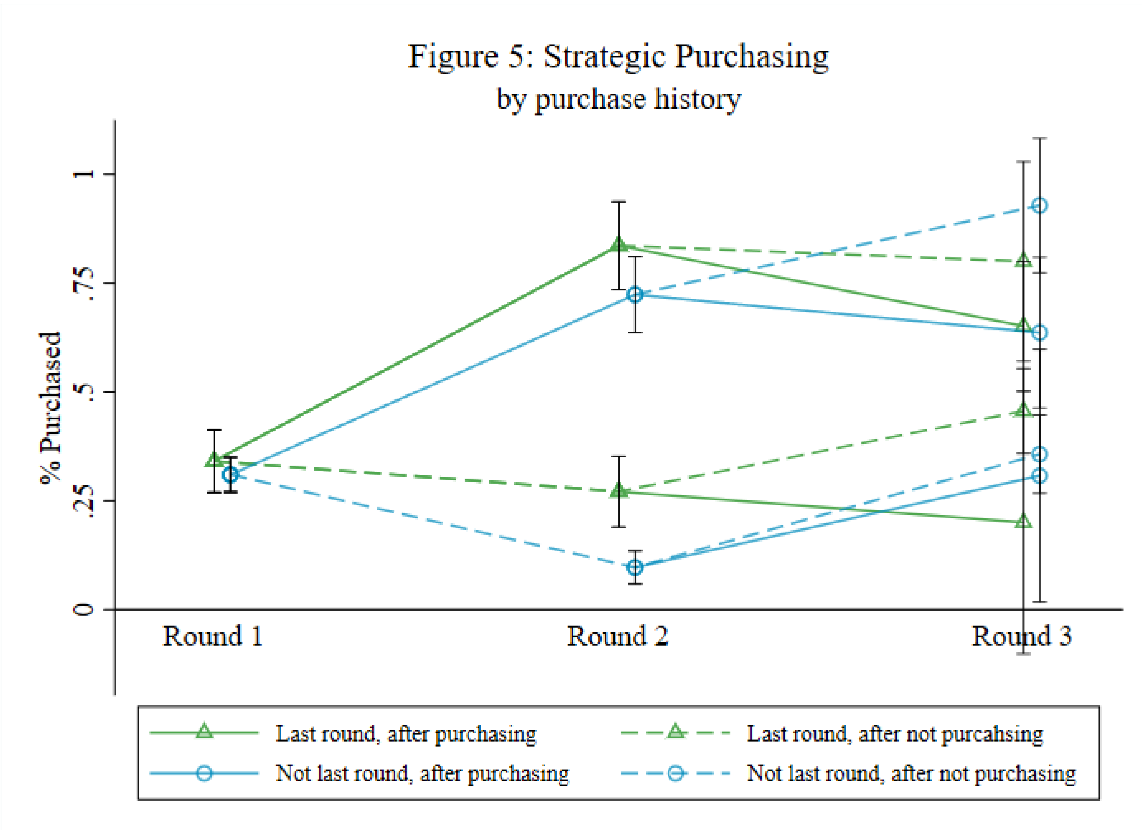

The results presented in Figures 4 and 5 indicate that participants strategically avoided purchases so as to be offered lower prices by the algorithm in subsequent rounds. However, this finding is not informative about the social costs this behavior entails. Indeed, participants could be making fewer purchases but where the baseline level (i.e., if purchasing did not trigger price changes) is socially excessive, or they could be forfeiting *efficient* transactions to signal their WTP. As indicated in Table 1, we found some evidence of participants avoiding efficient transactions in anticipation of future prices in treatment groups 1R–4R (H2), which we will now turn to present, and much clearer evidence in the information disclosure (D) condition (H4), which will be elaborated on in Section 5.

Figure 6, below, plots the percentage of participants who did *not* purchase the gift card *and* their WTP was equal to or higher than the price offered to them in the given round. In this figure, the percentage of socially inefficient avoidance of purchases is plotted on the y-axis,[14] which, by construction, is more statistically difficult to identify, as it does not make use of participants whose WTP was lower than the offered price but might have avoided the purchase for strategic reasons as well if the price was lower (or their WTP higher). Figure 7 takes the further step of altogether excluding participants whose WTP was lower than the offered price, sacrificing statistical power to arrive at an easily interpretable measure that captures the percentage of inefficient purchase avoidances out of all participants offered a seemingly beneficial price. The divergences between participants playing their final round and those expecting to play additional rounds were in the expected direction for all rounds, but were statistically significant at the 95% level only in round 3, thus somewhat supporting Hypothesis 2, but less than had been expected.

Figure 6: Strategic Purchasing

avoiding efficient purchases

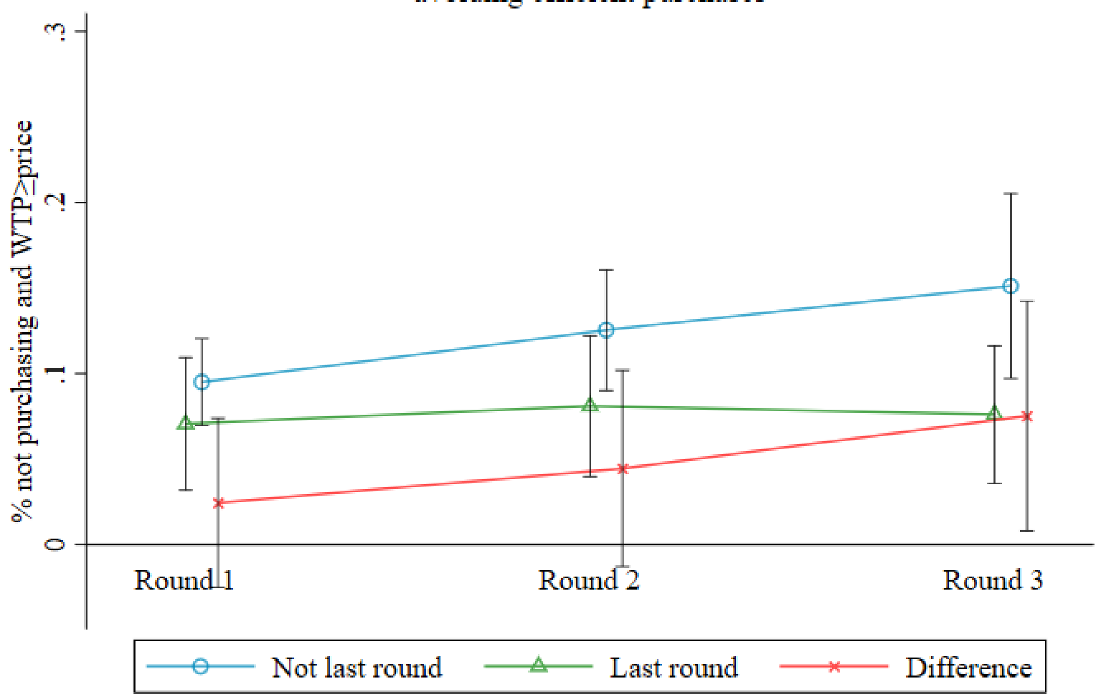

[14] As mentioned before, this abstracts away from the cost of production. If prices are higher than costs, which is usually the case with price discrimination, then social harm may be higher, as avoiding a purchase when WTP is lower than the price but potentially higher than the cost of production is a social waste.

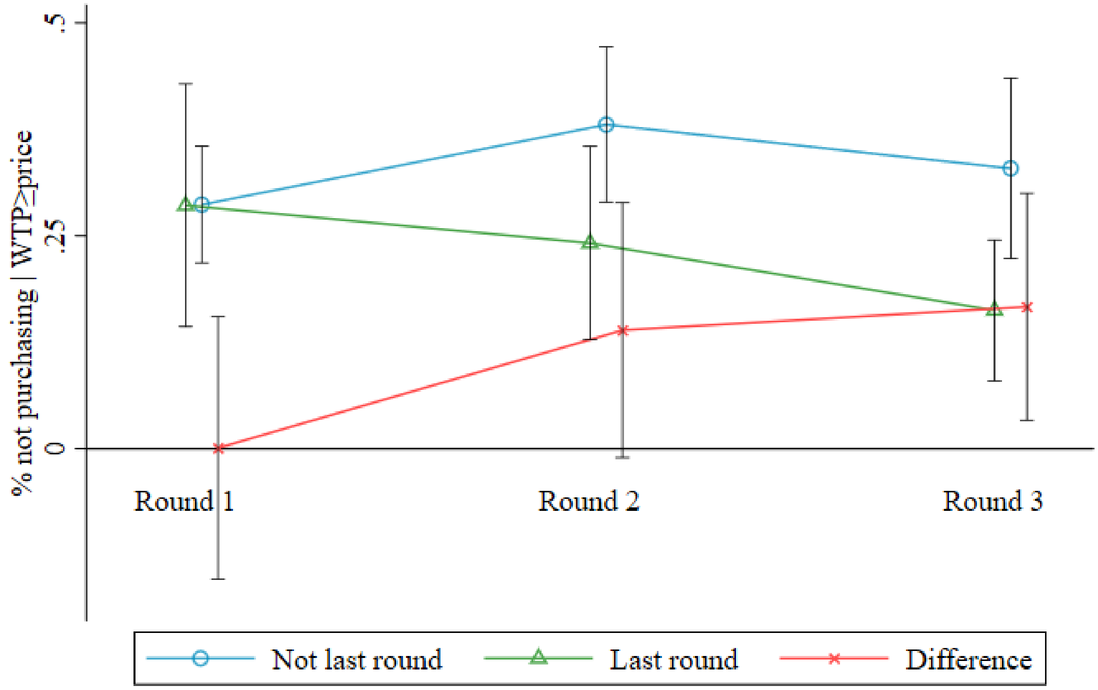

Figure 7: Strategic Purchasing
avoiding efficient purchases conditional on WTP≥price

5. Regulating Algorithmic Price Discrimination

5.1 Information Disclosure

Mandating an information disclosure is the most common tool used to regulate consumer markets. We simulated this mandate for participants in treatment groups C, D and E by presenting them with the following information in bold font at the beginning of the experiment, after the main comprehension test: "Note that your decision whether to purchase or not to purchase the gift card in any given round at the given price might affect the price that the algorithm will offer you in the next rounds – either increase or decrease it." As discussed in the Introduction, the effectiveness of disclosure mandates has been widely criticized on two central grounds: one, that consumers do not read them, and two, even if they do read them, they do not know how to process the information. Our study abstracted away from the first challenge by using an unrealistically short and simple disclosure that participants were forced to at least observe before proceeding to the next screen. Indeed, we found that the median participant spent 23.9 seconds reading the disclosure and that only 10% of the participants spent less than 10 seconds on the disclosure screen. We further found that almost all of the participants (96.35%) demonstrated a minimal level of comprehension of the information by correctly answering a test question presented on the next screen that read, "[I]f you choose to either buy or not buy the gift card in Round 2, could this affect the price that the algorithm will charge for the gift card in Round 3?"[15] In sum, the results presented below in this section are in no way intended as evidence of whether or not consumers read information disclosures but, rather, whether and how they are expected to use the information, if it is effectively conveyed to them, to improve their decision-making when engaging with algorithmic price discrimination.

Figures 8–11 replicate Figures 4–7 presented above, comparing participants' behavior in the Disclosure condition (in purple) to participants playing their final round without disclosure. Note that we can now plot round 4 in the figures, in contrast to the figures in Section 4, where there was only one condition (4R) in which participants played four rounds.

For each round, Figure 8 plots the percentage of participants who purchased the gift card when it was their final round without disclosure (in green), with disclosure (in purple), and the difference between the two groups (in red). As was the case without disclosure, we did not find any statistically significant divergences in participants' behavior in the first round, implying that the information alone was not sufficient to trigger strategic behavior before experiencing a price adjustment. However, we found that in rounds 2 and 3, participants were significantly less likely to purchase the gift card under the disclosure condition ($p < 0.01$ and $p = 0.039$, respectively). In contrast, without disclosure, we found a statistically significant effect on behavior only in round 2, and moreover, with disclosure, the effect was much greater in both rounds that it presented: 20.5 percentage points fewer purchases in round 2 and 11.1 percentage points fewer purchases in round 3. As expected, participants' purchasing decisions in round 4 were indistinguishable with and without disclosure, as this was the final round for all participants who reached it, and there were no longer strategic incentives to avoid purchases. These results provide clear evidence in support of Hypothesis 3, namely, that disclosure increase consumers' strategic avoidance of purchases. Furthermore, these results support an incentives tradeoff in the interplay between the participants and the pricing algorithm, where, on the one hand, participants learn to behave strategically, over

[15] However, 23.22% of participants had to return to the preceding screen with the information disclosure before responding to the comprehension question.

time, through their experience with the pricing algorithm but, on the other hand, their strategic incentives diminish as the number of rounds remaining decreases. This yields the concave shape of the red differences curve in Figure 8: inclining with the increase in experience and then declining with the decrease in strategic incentives.

Figure 8: Strategic Purchasing with Disclosure

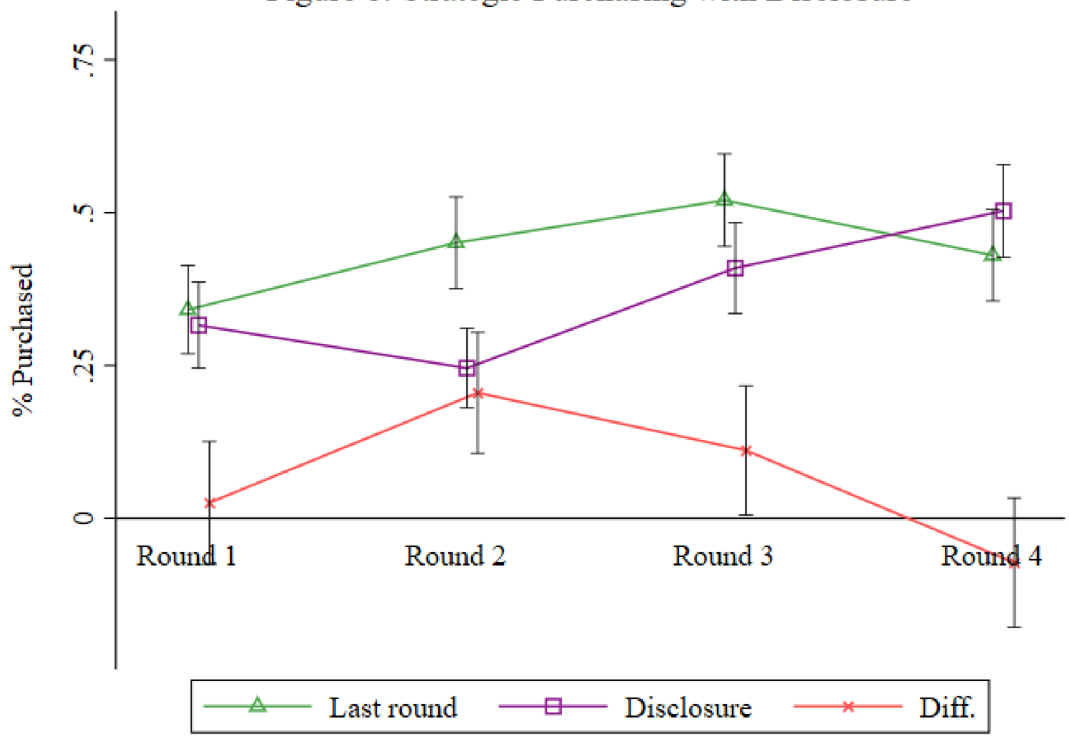

Figure 9, below, plots the percentage of purchases of participants in the Disclosure condition for every possible purchase history (and, accordingly, price), where a dashed line represents participants who did not purchase the gift card in the preceding round, and a solid line represents those who did purchase the gift card in the preceding round. The main results (presented in Figure 8) are robust to analyzing separately groups of participants who made the same decisions and were offered the same prices in all previous rounds: in round 2, participants with disclosure were significantly less likely to purchase the gift card regardless of whether they had purchased it in the previous round and were then offered a price of \$8 ($p < 0.01$) or had not made the purchase in the previous round and were then offered a price of \$6 ($p < 0.01$). Again, there was not enough statistical power to analyze participants' behavior in round 3 as there were four distinct purchasing histories that could lead to that round, thus diluting the sizes of some of the groups to 10 and 11.[16]

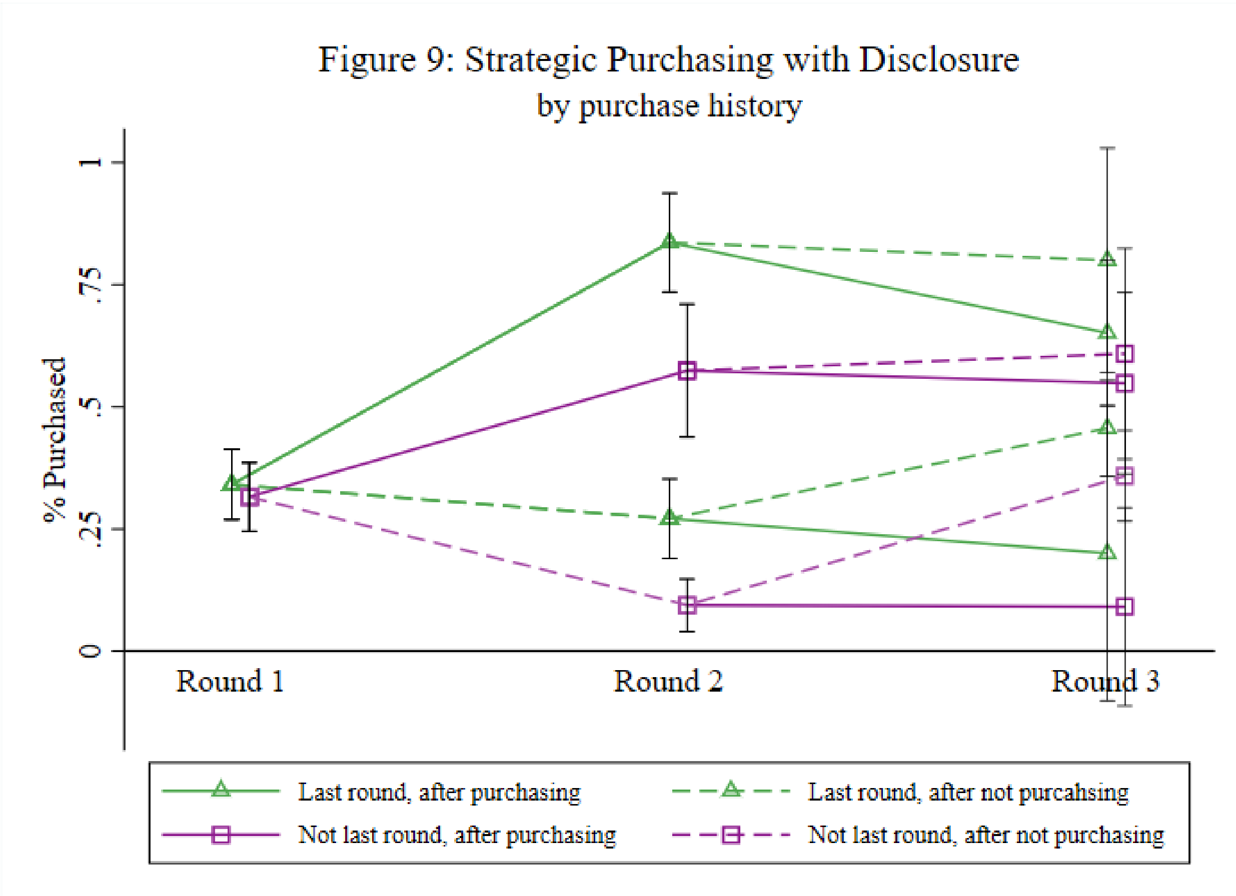

Finally, Figures 10 and 11 present a deeper dive into whether and to what extent participants in the Disclosure condition strategically avoided efficient transactions, i.e., when their WTP was equal to or higher than the offered price. Whereas in Section 4 we found only some evidence that participants strategically avoided efficient transactions in the absence of disclosure, this effect was clearer and stronger with disclosure, dramatically greater both in size and statistical significance. Furthermore, in both Figure 10 and Figure 11, there is the same concave pattern that presents in Figure 8, indicating that participants became more strategic as they engaged with and experienced the pricing algorithm and then gradually less strategic in each round as incentives

[16] Surprisingly, despite the lack of statistical power for round 3, for each possible purchase history, the divergence in purchasing decisions of participants in round 3 with and without disclosure are consistent, despite not being statistically significant, with strategic avoidance of purchases. This can be observed by the fact that for each pair of green triangle and purple square in round 3 that originate from the same point in round 2 and share the same line pattern (i.e., dashed or solid), the triangle lies above the square.

diminished, dropping off completely in round 4 when there was no longer any strategic reason to avoid an efficient transaction. Especially telling are the results presented in Figure 11, showing that with disclosure, a whopping 55% of the participants who were offered a price equal to or lower than their WTP declined the purchase in round 2 and 42.4% declined in round 3. This is in contrast to only 24.1% and 16.3%, respectively for rounds 2 and 3, for participants playing their final round without disclosure (the differences between the two groups in both rounds are significant at $p < 0.01$).

Figure 10: Strategic Purchasing with Disclosure

avoiding efficient purchases

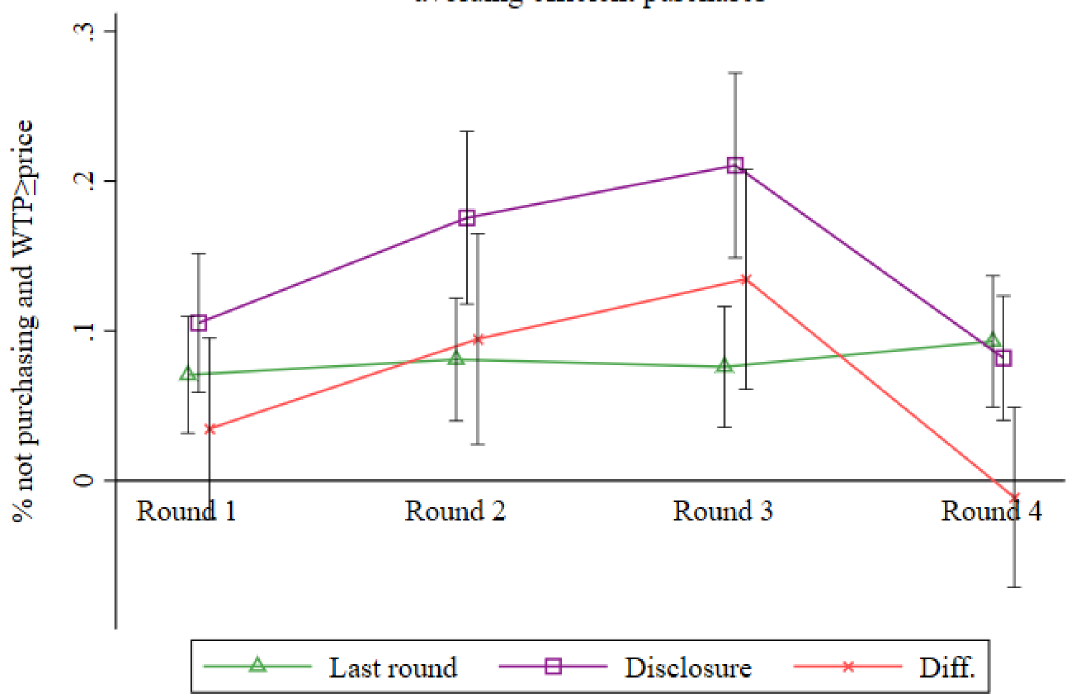

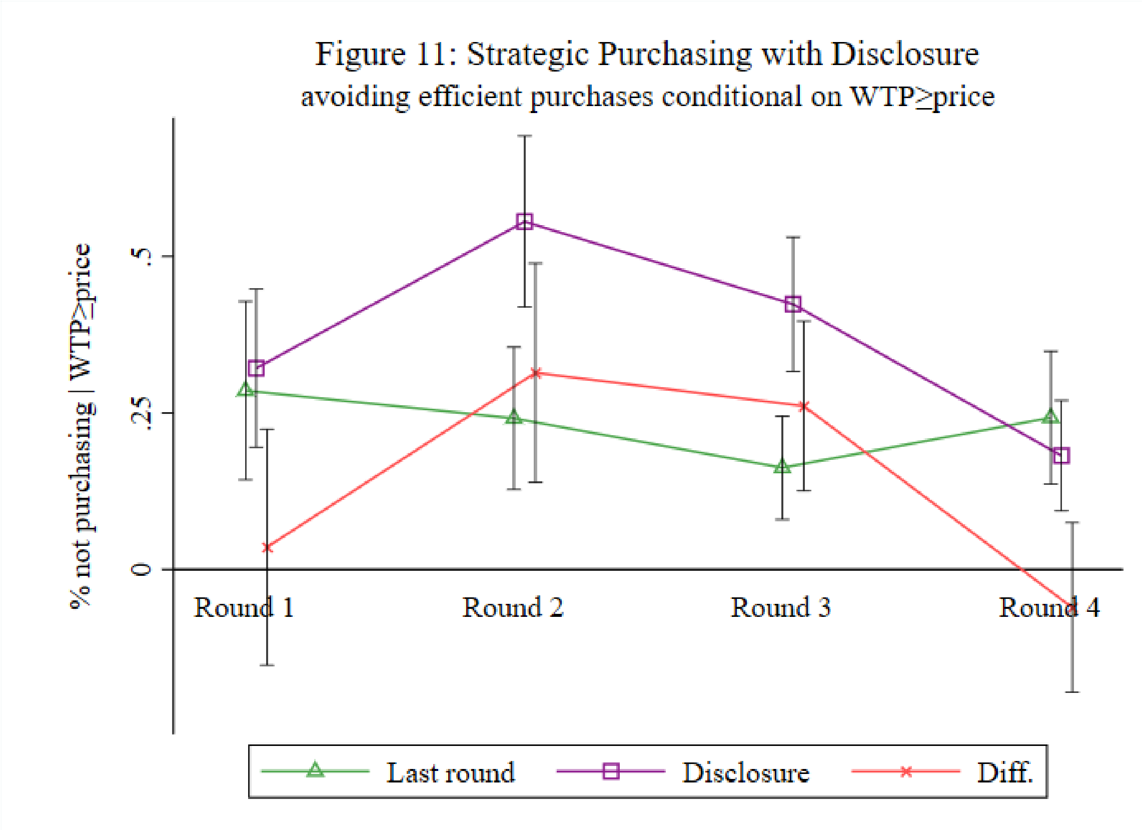

5.2 Data Protection

The second framework for regulating algorithmic price discrimination investigated in our study is data protection rights. Because algorithmic pricing uses information gleaned from consumers' behavior, giving consumers control over the data collected on them in online markets is a straightforward mechanism for regulating algorithmic price discrimination. Specifically, we simulated the informational functions of two of the most prominent (yet under-researched) data protection rights: the right to prevent data collection ex ante (most often in the form of a right to refuse the collection of cookies) and the right to delete data ex post (the right to erasure or "to be forgotten"). The many differences between these two rights could potentially render either one of them more effective for regulating algorithmic pricing. However, to rigorously identify a causal mechanism through which these rights may have a different effect on consumer behavior, we simulated them in our study so that they only differed in the timing of the presentation of the choice to allow data collection: before participants made their purchase decision for the Cookies condition (but after learning about the price for that round) and after the purchase decision for the Erasure condition. Accordingly, we intentionally formulated the cookies and erasure conditions identically, thus abstracting away from differences in the specific formulations of each right in the real world to avoid confounding and obscuring the causal source of any observed effects. Specifically, in both conditions, participants were informed, “Before [making your choice *or* proceeding to the next round], please indicate whether you would like the algorithm to use the information about your decision to set prices in future rounds.” Importantly, participants were offered these choices only from round 2, thus ensuring that all experienced a price change at least once, since there was no price change when participants opted to prevent data collection. Second, recall that participants in these two conditions were also given the information disclosed to group D, thereby allowing us to identify the added effects on consumer behavior of these rights beyond the effect of disclosure.

Admittedly, these rights were formulated and presented in a way that does not purport to emulate real-world data protection rights, much as the information disclosure did not resemble real-world disclosures: the rights were phrased concisely, their existence was salient to participants (rather than hidden in fine print), their potential effect on future prices was clear, and utilizing them was costless. Therefore, the findings presented in this section should not be interpreted as implying that real-world consumers use data protection rights in the same way as participants did. In fact, it is probably safe to assume that most consumers are oblivious to the informational effects of data protection rights that this study investigates. Instead, the purpose of this study is to shed light on the potential interplay between data protection rights and strategic consumer behavior, which, in turn, should inform how these rights are structured, regulated, and enforced. In other words, rather than constraining research by the ways in which data protection rights are currently exercised by firms and consumers, policymaking should be informed by the potential informational effects of data protection rights when deciding how to regulate issues such as the disclosure of these rights, the way they are phrased, the hurdles that firms often create to prevent consumers from using them, and so on.

Figure 12 plots, for each round and for each of the two data protection conditions, the percentage of participants who decided to allow data collection. We found that a large majority of participants allowed data collection (around 79.5% on average, across rounds and conditions). This was not surprising given the high rates of consumer consent to cookie collecting found in observational studies, a tendency that the acquiescence bias likely intensifies in a lab setting. More importantly, however, these high data collection rates do not imply that participants were unsophisticated. Across all rounds and experimental conditions, 63.1% of the participants declined the purchase, in which event the strategic participant would have benefited from allowing data collection (had they been given the choice), as this would lead to a price decrease in the next round, rather than disallowing data collection, which would result in no change to the current price. Figure 12 indicates that there was no statistically significant divergences between the Cookies and Erasure conditions in data collection rates. While data collection was slightly lower for the Erasure treatment group in all rounds, which is consistent with the logic motivating Hypothesis 5 – that making the choice ex post is more conducive to a contemplative consideration of the legal right – the lack of statistical significance to this divergence means that no evidence was found to support Hypothesis 5 using this test. However, some of the other results that will be presented later in this section do provide some evidence of more nuanced divergences in the data collection and retention patterns between the two groups.

Figure 12: Data Retention with Cookies and Erasure

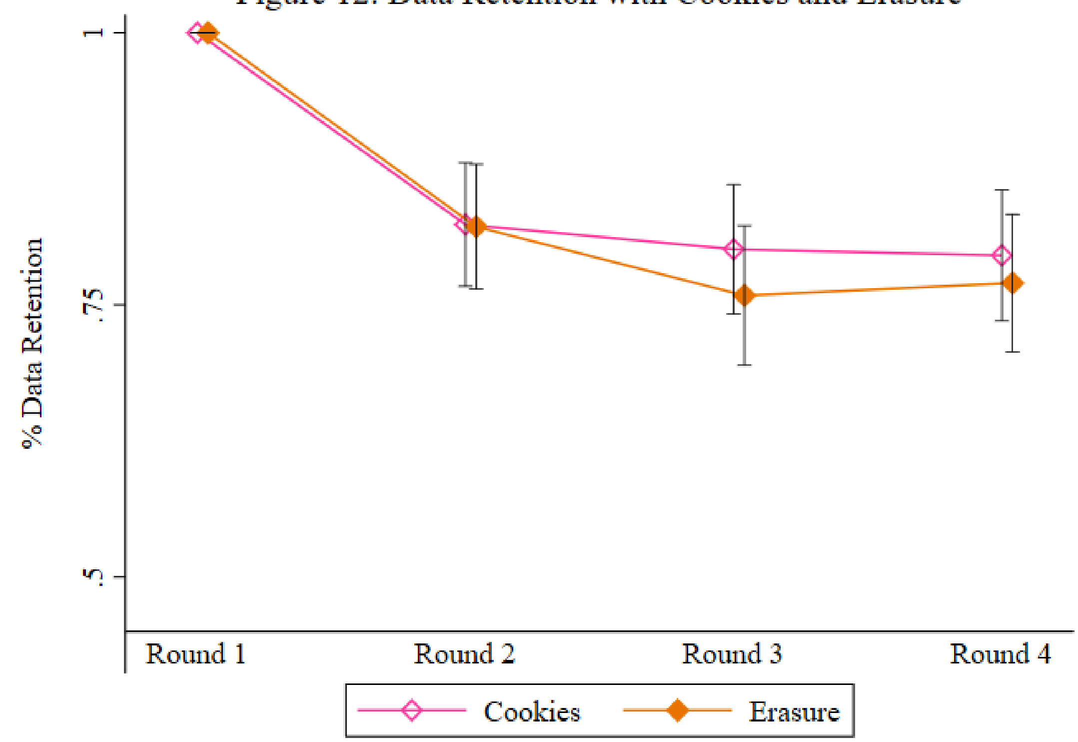

When given control over their data, participants were expected to be more willing to purchase the gift card if the price was lower than their WTP because they could avoid a price increase in the next round by refusing data collection. However, by so doing they would not enjoy the price reduction that would result from avoiding the purchase but allowing data collection. Therefore, there was incentive for these participants to strategically avoid the purchase and signal a lower than their actual WTP to the pricing algorithm by allowing data about their purchase decision to be retained, but it was weaker when compared to the strategic incentives of participants the other experimental condition that had no control over data collection. Figure 13 shows that participants with control over their data tended to behave strategically in round 2, with their average purchase rate lower than participants in the 2R condition who were now playing their final round but indistinguishable from the purchase rate of participants in the Disclosure condition. However, it is impossible to draw any inferences about the potential effects of the data protection rights in rounds 3 and 4, as the option to prevent data collection from round 2 onward meant that participants in the Cookies and Erasure conditions faced different prices in rounds 3 and 4 than participants in the Disclosure condition, undermining a central identification assumption that, in each round, the groups of participants being compared had faced similar prices (on average) in the rounds leading up to it.

Figure 13: Strategic Purchasing with Cookies and Erasure

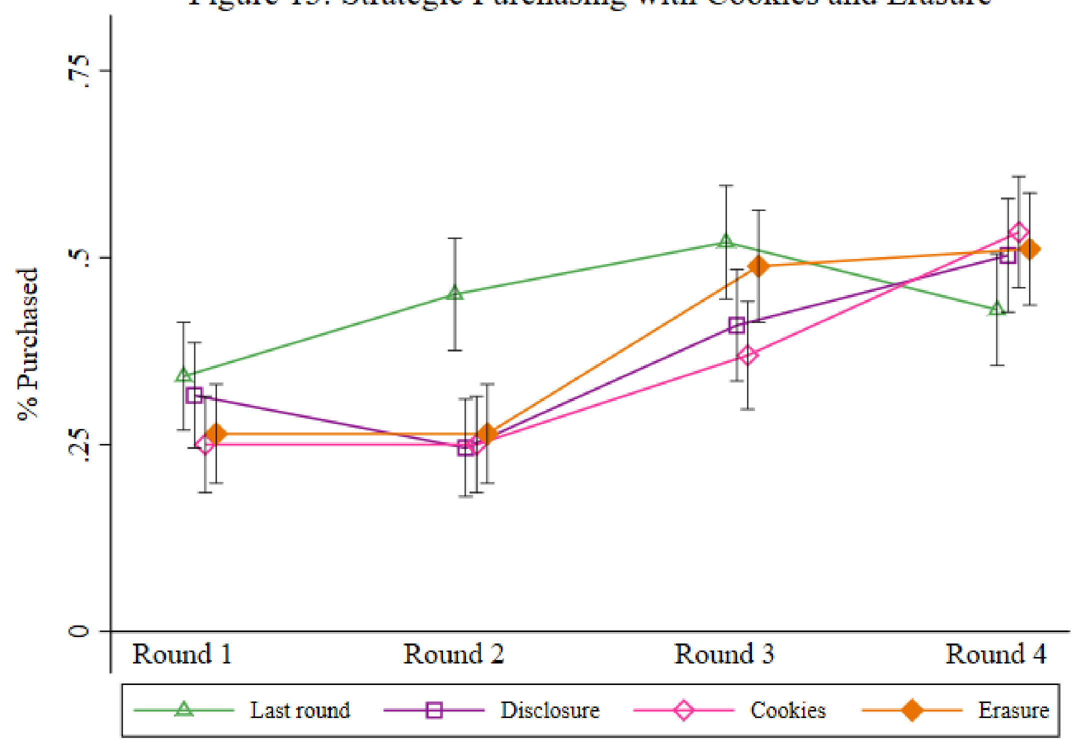

To address this limitation and offer an answer to whether participants in the Cookies and Erasure conditions used their data protection rights strategically, we first plotted, in Figures 14 and 15, the purchase rates of the two groups separately for all data retention histories. Note that in this figure, as opposed to similar analyses presented in Figures 5 and 9, a dashed line leading to a specific round indicates that data was not collected in the round where the line *ends*, whereas a solid line indicates that participants allowed data collection. We found that for both conditions, the purchase rate of participants who allowed data retention in round 2 was indistinguishable from the Disclosure condition, while the purchase rate of participants who refused data collection significantly higher, offering some evidence in support of Hypothesis 7, that data protection rights decrease strategic avoidance of purchases (but again, there was no statistical power to analyze rounds 3 and 4 in this approach for controlling for different data collection histories). However, a caveat is in order: these results provide only suggestive evidence of the hypothesized effect because of the possibility that, for whatever reason, participants who allowed data collection in the second round exhibited different purchasing patterns in the first round compared to participants who prevented data collection. This would have led to different prices being offered to these two groups in the second round, which would affect their purchasing decision in a way that is unrelated to the strategic reason to either allow or prevent data collection.

Figure 14: Strategic Purchasing with Cookies
by data retention history

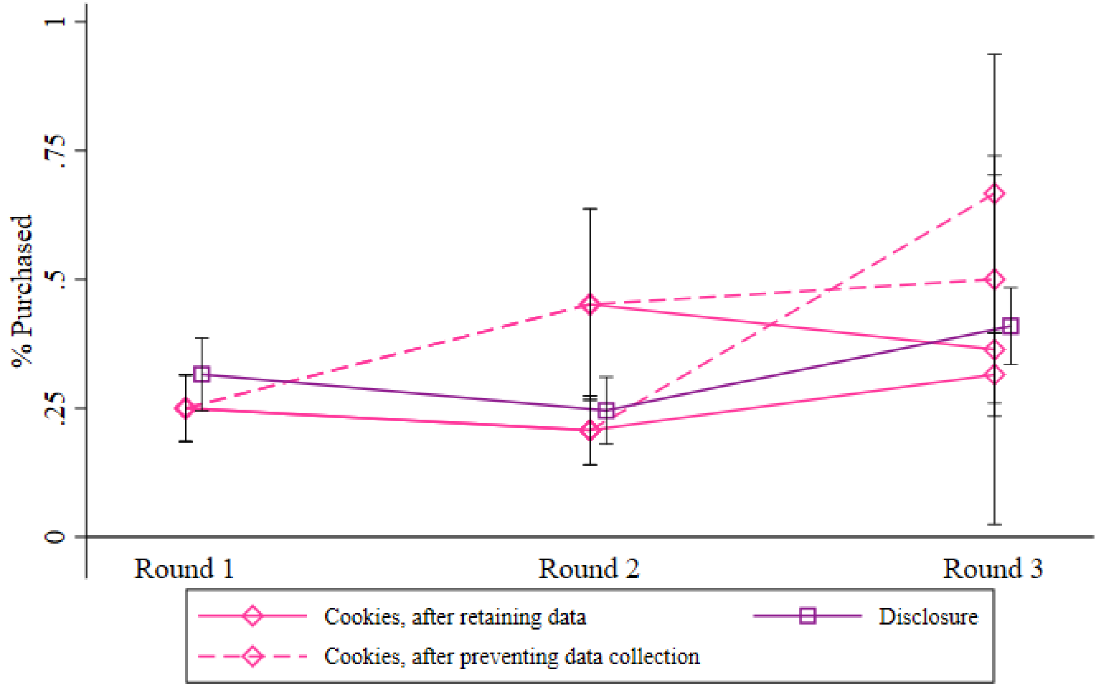

Figure 15: Strategic Purchasing with Erasure
by data retention history

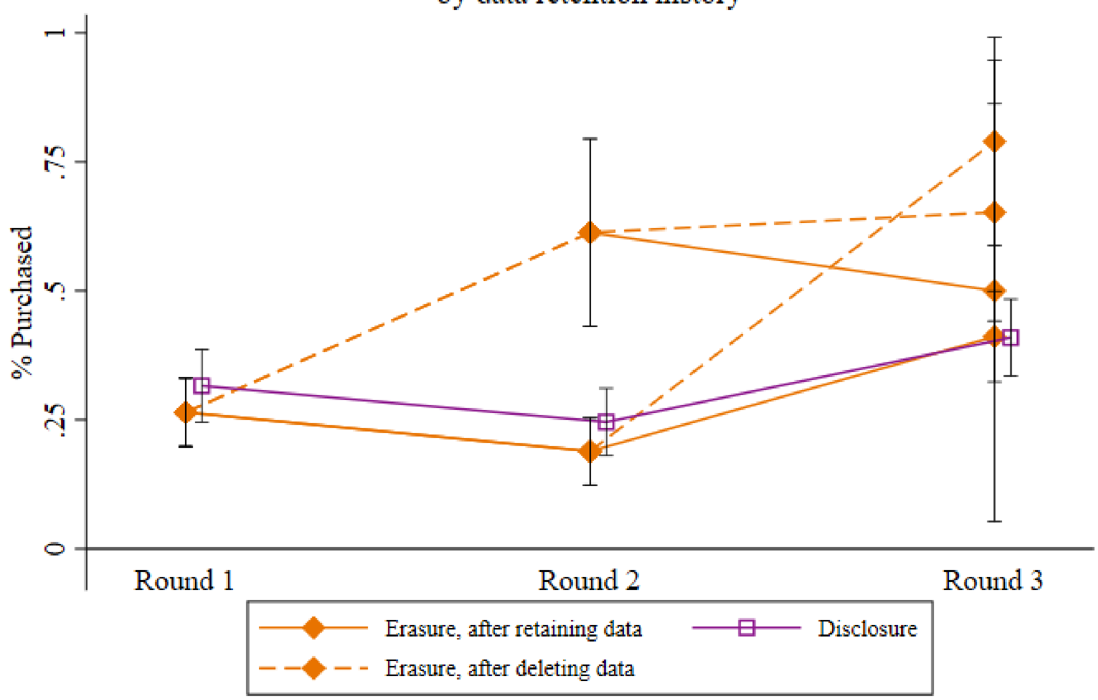

Finally, we derived the most substantial and robust evidence that participants in the two conditions used their data protection rights strategically, as presented in Figures 16 and 17, by comparing the data collection rates *by purchase decisions*, for each round and experimental condition. As opposed to prior analyses of participants' strategic use of their data protection rights, this analysis has the important benefit of relying on a stark theoretical benchmark for the strategic use of these data protection rights: *all* participants who decided not to purchase the gift card in any given round could rationally benefit from strategically allowing data collection, signaling a low WTP to the pricing algorithm and thus gaining a lower price in the next round; at the same time, *no* participant who purchases the gift card should rationally allow data collection, as that would lead to a price increase when preventing data collection would preserve the price at its current level. As to be expected, participants' behavior significantly deviated from this perfectly rational benchmark. More important to our purposes is that we found that, on average, participants were significantly more likely to allow data collection when they did not purchase the gift card ($p < 0.01$ for both conditions and both rounds 2 and 3), effectively signaling a low WTP to the pricing algorithm, which provides clear evidence for Hypothesis 6, namely, that consumers are more likely to allow data collection when not purchasing.[17] As an aside, it seems that the rate of strategic data collection may be greater under the Erasure condition than the Cookies condition. While these differences-in-differences across the conditions are not statistically significant at the 95% level, they may constitute some suggestive evidence in support of Hypothesis 5,[18] consistent with the theoretical conjecture that an ex post data collection decision is more contemplative and, thus, more strategic than an ex ante decision.

[17] We present the results for round 4 for full transparency but refrain from drawing any inferences from participants' behavior in that round: since data collection or retention in the final round is inconsequential, the rational participant would be indifferent between allowing or preventing her data from being collected. Moreover, presenting this choice in round 4 may have caused confusion by implying to participants that their decision could somehow still affect a future bonus payment. In these circumstances, it is to be expected that participants may exhibit a preference for either allowing or preventing data collection in round 4, but it would be imprudent to infer anything meaningful from such behavior.
[18] Namely, that consumers make different use of data protection rights when invoked before versus after a transaction.

Figure 16: Data retention as function of purchase decisions
with a right to prevent data collection (cookies)

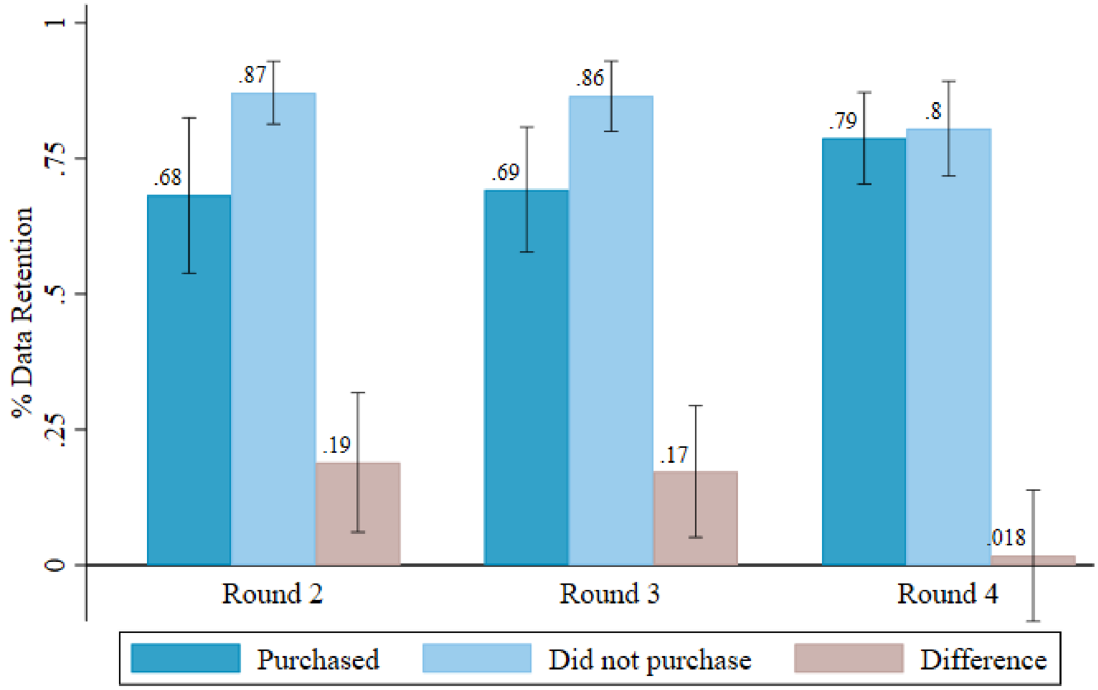

Figure 17: Data retention as function of purchase decisions
with a right to delete data (erasure)

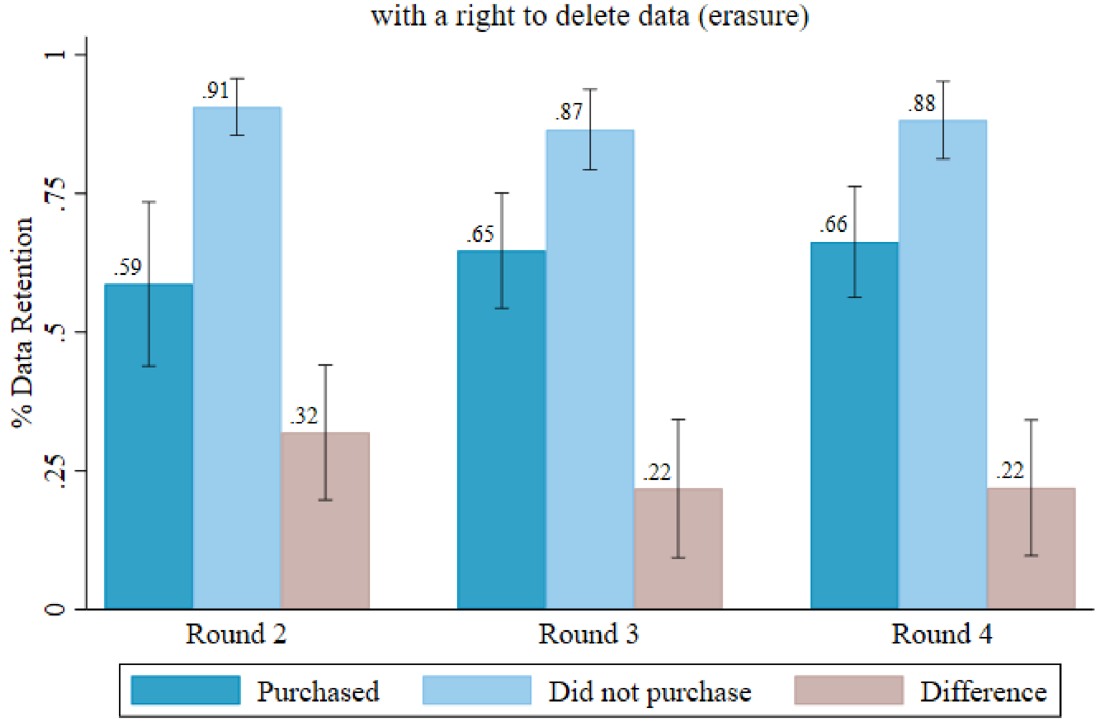

Before concluding, it could arguably be asserted that even these results may be driven by factors other than strategic use of data protection rights. For example, individuals may have some unobserved personal characteristic that renders them more inclined, for whatever reason, to purchase the gift card and prevent data collection and not because their purchase decision motivated their data collection decision. For this reason, we conducted a final set of robustness tests for Hypothesis 6 in the form of a series of regressions of the decision to allow data collection on the decision to purchase the gift card, with the results reported in Table 6 below. Column 1 shows the results of a simple regression that roughly replicates the results presented in Figures 16 and 17, indicating that the decision to purchase is negatively associated with the decision to allow data collection. Crucially, the regression models whose results are presented in columns 2 to 5 utilize the multi-round design of the experiment to estimate the effect of the purchasing decision on the data collection decision by including individual fixed effects in the estimation, harnessing the variance in participants' decisions (on whether to purchase the gift card and whether to allow data collection) across rounds. These fixed-effects models allowed us to identify the within-person effect of the purchasing decision on the decision to allow data collection, controlling for all potential unobserved personal attributes that may contribute to the observed negative association. Indeed, model 2, which incorporates individual fixed effects, shows a decrease in the magnitude of the negative association between the purchase and data-collection decisions, but the effect is still negative, large, and significant ($p = 0.011$). Model 3 excludes round 4, where data retention is inconsequential from a rational agent's perspective, which, as expected, increased the estimated effect's magnitude and significance. Finally, models 4 and 5 are specified similarly to model 3 but estimate the effect separately for the Cookies and Erasure conditions. They present the interesting result that, with individual fixed effects, both coefficients are negative, but only the one obtained for the Erasure condition is statistically significant. This provides further suggestive evidence in support of Hypothesis 5, namely, that providing consumers with a right that allows ex post control over their data might be utilized differently (and more effectively) than a right that grants ex ante control over data, i.e., before a purchase decision is made.

Table 6: Strategic Data Collection and Retention

	(1)	(2)	(3)	(4)	(5)
Purchased	$-.16^{***}$	$-.073^{*}$	$-.14^{**}$	-.074	$-.2^{**}$

	(.025)	(.029)	(.045)	(.065)	(.062)
Price	-.05***	-.0019	-.0083	-.02	.0049
	(.0083)	(.017)	(.029)	(.04)	(.042)
Constant	1.2***	.84***	.9***	.96***	.83**
	(.051)	(.11)	(.19)	(.26)	(.28)
Fixed effects	No	Yes	Yes	Yes	Yes
Excluding round 4	No	No	Yes	Yes	Yes
Including Cookies	Yes	Yes	Yes	Yes	No
Including Erasure	Yes	Yes	Yes	No	Yes
Observations	1050	1050	700	352	348
Individuals	350	350	350	176	174
Adjusted R^2	.08	.58	.54	.51	.57

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

6. Discussion and Concluding Remarks

Algorithms have had a dramatic impact on many economic practices, including enabling sellers to set highly granular personalized prices based on algorithmic predictions of consumers' WTP. In the algorithm era, WTP predictions rely primarily on consumers' behavior rather than immutable characteristics. However, the estimations of even the most sophisticated algorithms cannot be relied on without knowing whether consumers' behavior, which is presumed to signal their WTP, represents their true preferences or whether some consumers might adjust their behavior, at least partially, to send signals that will result in better transaction outcomes for them. This paper takes an initial swing at empirically testing the plausibility of the latter possibility, namely, that consumers may behave strategically in engaging with pricing algorithms to effectively "bargain" for better prices, which has crucial implications for the theoretical analysis of algorithmic behavior-based price discrimination and its regulation.

Note that no normative stance is taken in the paper on whether personalized algorithmic price-discrimination and its regulation are good or bad (for consumers or the market as a whole), a daunting task undertaken elsewhere (Porat, 2022). Rather, we seek only to empirically demonstrate that consumers can reclaim some of their long-lost market bargaining power when faced with a pricing algorithm and that many are sufficiently sophisticated to do so, especially when in possession of sufficient information, experience, and control over their data. For decades, sellers used uniform, cold-hearted boilerplate contracts to dominate transactions and essentially dictate the transaction terms, rendering consumers powerless. Ironically, the non-human, artificially intelligent algorithm—the technologically advanced price-setting tool that has supposedly enhanced sellers' sophistication and advantage in the contractual relations—has resurrected human bargaining dynamics in the consumer market. In addition to empirically demonstrating consumers' potential capacity for taking strategic advantage of the advent of algorithmic price-setting, this paper illuminates the potential of regulating the practice, shedding novel functionalist light on overlooked aspects of consumer and data protection laws that can be applied beyond consumers' privacy (and, likely, beyond algorithmic pricing).

Finally, to be clear, we acknowledge that the finding that participants in laboratory experiment learn to behave strategically in interacting with a price-setting algorithm cannot be assumed to perfectly predict how consumers behave in the real world, which is for future observational research to explore. In the real world, consumers rarely observe how their behavior affects prices directly, as prices may be affected by general market conditions; the time gap between consecutive transactions may be long (and the expected number of transactions may be either lower or much higher than four); and behavior might take various forms, each contributing to the price-setting function of the algorithm in obscure ways (e.g., browsing history, purchasing history of various products sold by the same seller, etc.). Furthermore, even in contexts where consumers can be reasonably expected to develop the technological literacy and savviness necessary to successfully negotiate with algorithms, giving them information or control over their data will unlikely level the playing field entirely. These protections may be ineffective or insufficient in many concrete contexts (disclosures are lengthy and obscure, and exerting control over one's data is time-consuming). As with all experimental research, this study is not driven by an ambition to capture a 4K-pixel picture of the world but rather to sketch an accurate architectural design of its fundamental building blocks, i.e., to offer focused insight into the incentives and behavioral mechanisms underlying the strategic interplay between consumers and pricing algorithms when different forms of regulation are implemented, which will hopefully help refine the assumptions underlying the theoretical study of algorithmic pricing and inform more nuanced policymaking when it comes to regulating algorithmic pricing using consumer and data protection tools.

References

Acquisti, A., Taylor, C., & Wagman, L. (2016). The Economics of Privacy. *Journal of Economic Literature, 54*(2), 442-492.

Bakos, Y., Marotta-Wurgler, F., & Trossen, D. R. (2014). Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts. *The Journal of Legal Studies, 43*, 1-35.

Bar-Gill, O. (2019). Algorithmic Price Discrimination: When Demand Is a Function of Both Preferences and (Mis)Perceptions. *The University of Chicago Law Review, 86*(2), 217-254.

Bar-Gill, O., Sunstein, C. R., & Talgam-Cohen, I. (2023). Algorithmic Harm in Consumer Markets. *Available on SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4321763*.

Becker, G. M., Degroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science, 9*(3), 226–232.

Ben-Shahar, O., & Schneider, C. E. (2014). *More Than You Wanted to Know: The Failure of Mandated Disclosure.* Princeton University Press.

Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive Consumer Choice Processes. *Journal of Consumer Research, 25*(3), 187–217.

Chen, L., Bó, I., & Hakimov, R. (2023). Strategic Responses to Personalized Pricing and Demand for Privacy: An Experiment. *SSRN*. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4562169

Gillis, T. B., & Spiess, J. L. (2019). Big Data and Discrimination. *The University of Chicago Law Review, 86*, 459-487.

Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. *Journal of Legal Analysis, 10*, 113-174.

Marotta-Wurgler, F. (2016). Self-Regulation and Competition in Privacy Policies. *The Journal of Legal Studies, 45*, 513-539.

Marotta-Wurgler, F., & Davis, K. E. (forthcoming, 2024). Filling the Void: How E.U. Privacy Law Spills Over to the U.S. *Journal of Law and Empirical Analysis*.

Porat, H. (2022). Behavior-Based Price Discrimination and Data Protection in the Age of Algorithms. *Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4254326*.

Porat, H. (2024). Algorithmic Personalized Pricing in the United States: A Legal Void. In F. Esposito, & M. Grochowski (Eds.), *Cambridge Handbook on Price Personalization and the Law.* Cambridge University Press.

Shiller, B. R. (2020). Approximating Purchase Propensities and Reservation Prices from Broad Consumer Tracking. *International Economic Review, 61*, 847-870.

Taylor, C. R. (2004). Consumer Privacy and the Market for Customer Information. *RAND Journal of Economics, 35*(4), 631-650.

Villas-Boas, J. M. (2004). Price Cycles in Markets with Customer Recognition. *RAND Journal of Economics, 35*, 486-501.

Yoo, C. S. (2022). The Overlooked Systemic Impact of the Right to Be Forgotten: Lessons from Adverse Selection, Moral Hazard, and Ban the Box. *University of Pennsylvania Law Review Online, 170*.

Appendix: Experimental Protocol

Start of Block: Consent Form

You are invited to participate in an experiment on decision-making. You must be a US national, fluent in English and at least 18 years of age to participate. Your participation will take about 8 minutes, but once you begin you are advised to finish the survey. There are no risks associated with this study, and your identity will be kept confidential. As part of this research design, you may not be told about the purpose of this research, but all the information we do provide is completely accurate.

Participation: If you decide to participate in this study, please note that your participation is voluntary and that you may withdraw your consent or discontinue participation at any time without penalty. Your privacy will be maintained in all published and written data resulting from the study. Your name will never be connected to any decision you make.

IMPORTANT: you may only take this survey **ONCE**. If you exit the survey midway and restart it, for whatever reason, your submission, if not returned, will be rejected, both preventing you from being paid and adversely affecting your Prolific approval rate.

Payment: In addition to the payment for completing this survey, you will have a chance to receive a bonus payment, based partly on your decisions, and partly on chance, as detailed in the instructions that follow.

Contact Information: If you have any questions, concerns, or complaints about this research, its procedures, risks and benefits, contact hporat@sjd.law.harvard.edu.

By clicking "Agree" below, you confirm that you have read the consent form, are at least 18 years old, fluent in English, a US citizen, and agree to participate in the research.

- ◯ I agree to participate
- ◯ I do not agree to participate

Before proceeding, please help us ensure that you are not a robot:

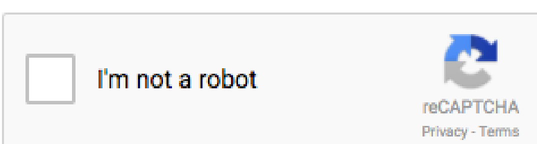

End of Block: Consent Form

Start of Block: Prolific ID

What is your Prolific ID? Please note that this response should auto-fill with the correct ID

____

End of Block: Prolific ID

Start of Block: instructions

Instructions

In this study, you will have a chance to win a bonus payment of $10 (or potentially more), in addition to the participation fee, by making several everyday purchasing decisions. Specifically, a random 10% of participants will be awarded a bonus payment, which will be transferred within a week after the completion of the survey. The size of the bonus payment will depend on the decisions that you make in this survey, as detailed below. All participants have an equal chance of winning the bonus payment, and the *likelihood* of winning the bonus has nothing to do with your decisions or the bonus size.

This study consists of up to 4 rounds, randomly determined by the computer. You will be informed whether you will play 1, 2, 3 or 4 rounds only before the beginning of the last round. For example, if you are chosen to play 3 rounds, then you will be told that round 3 is the last round before starting it.

In each round, you will first be given $10 as your potential bonus. You will then be asked whether you want the bonus for that round to consist of the entire $10 in cash, or to use some of it to buy a Walmart gift card with a value of $10. For example, if you are offered a price of $7 for the Walmart gift card, you may choose to buy the gift card, in which case your bonus (for that round) will consist of a $10 Walmart gift card and $3 in cash (your change, after spending $7 on the gift card). Instead, you may choose not to buy the gift card, in which case your bonus (for that round) will be $10 in cash. The gift card can be used like cash in any Walmart store in the United States (as well as online) and never expires, and it will be transferred to you via the Prolific messaging interface in the form of a unique link. Your final bonus, if chosen to receive it, will be determined by the payoff you obtained in one of the rounds, randomly selected by the computer upon the completion of the study.

The price of the Walmart gift card in each round will be determined by an algorithm. **Importantly, the price may vary across participants and from one round to the next, and will depend on everything the algorithm learns about you.**

Specifically, the algorithm is programmed to set a personalized price for the gift card that will make it beneficial for you to choose the gift card over the cash payment and, as a result, shop at Walmart. However, the algorithm tries to set the highest price for which you will still choose the gift card (over the cash payment), much like a car dealer or a flea market vendor. For example, if the algorithm learns that you frequently shop at Walmart, it will set a higher price, as you would still choose the gift card over the cash payment. On the flip side, if the algorithm learns that you rarely shop at Walmart, it will set a lower price to increase the likelihood that you will purchase the gift card.

Final notes: This research is conducted by Harvard University researchers and funded by Harvard University. The outlined procedure is entirely accurate and will be strictly followed by the researchers. Failing to comply with any of our commitments to you will expose us to either legal or ethical liability that could invalidate this study. In the next screen, we will ask you a few questions to ascertain that you have read these instructions and understand the mechanisms of the study.

End of Block: instructions

Start of Block: instructions_comprehension

Please answer the following questions, which are meant to ascertain your comprehension of the instructions in the previous screen. Each question has only one correct answer, and **you must answer the questions correctly to be eligible to receive the bonus payment**. Please feel free to return to the previous screen to re-read the instructions. Once you submit your answers, you cannot change them, and exiting the survey after submitting your answers and reentering it, for whatever reason, will result in the rejection of your submission.

Where can you use the gift card that you will be able to purchase during this survey?

____

Will you be offered the gift card for the same price in every round?

○ Yes

○ No

If you are offered the gift card at a price of $6 and decide to buy it, how much cash will your bonus include in addition to the gift card?

○ $0

○ $1

○ $2

○ $3

○ $4

○ $5

○ $6

○ $7

○ $8

○ $9

○ $10

If you rarely shop at Walmart, will the price of the gift card offered to you be lower or higher than that offered to a person who frequently shops at Walmart? (assuming the algorithm knows this about you)

○ Lower

○ Higher

Once you feel confident with your answers, please proceed to the next screen.

End of Block: instructions_comprehension

Start of Block: treatment_instructions (displayed to conditions D, C, E)

Important Disclosure

Note that your decision whether to purchase or not to purchase the gift card in any given round at the given price might affect the price that the algorithm will offer you in the next rounds – either increase or decrease it.

Display This Question:
If treatment = erasure

However, after every round, beginning in round 2, you will be given a choice whether to allow the algorithm to use your choice to adjust the price of the gift card in future rounds.

Display This Question:
If treatment = cookies

However, before every round, beginning in round 2, you will be given a choice whether to allow the algorithm to use your choice to adjust the price of the gift card in future rounds.

In the next screen, we will ask you one final question to ascertain that you have read this disclosure and understand the mechanisms of the study.

End of Block: treatment_instructions

Start of Block: disclosure_comprehension

Please answer the following question, which is meant to ascertain your comprehension of the information in the previous screen. The question has only one correct answer. Please feel free to return to the previous screen to re-read the information. Once you submit your answer, you cannot change it.

If you choose to either buy or not buy the gift card in Round 2, could this affect the price that the algorithm will charge for the gift card in Round 3?

○ Yes

○ No

Once you feel confident with your answers, please proceed to the next screen to begin Round 1.

End of Block: disclosure_comprehension

Start of Block: R1_offer

<u>Round 1</u>

$10 has been added to your bonus payment for this round. Based on the information we have about you, in this round the algorithm is offering you the $10 Walmart gift card **for a price of $7**.

Please proceed to the next screen to decide whether you want to use some of your bonus payment to purchase the gift card.

End of Block: R1_offer

Start of Block: R1_choice

<u>Round 1</u>

Please choose **one** of the following options as your bonus payment for this round:

○ A $10 cash payment

◯ A $10 Walmart gift card and the remaining $3 (the change) in cash

 

Note that you will not be able to change this decision once you proceed to the next screen.

End of Block: R1_choice

Start of Block: R1-4_feedback

Your choice of bonus payment for this round has been recorded. Please proceed to the next screen to continue to the next round.

End of Block: R1-4_feedback

Start of Block: R2_offer

Round 2

$10 has been added to your bonus payment for this round. Based on the information we have about you, in this round the algorithm is offering you the $10 Walmart gift card **for a price of ${ price2}**.

Display This Question:
If treatment = cookies

Before making your choice, please indicate whether you would like the algorithm to use the information about your decision to set prices in future rounds.

◯ use information

◯ do not use information

Please proceed to the next screen to decide whether you want to use some of your bonus payment to purchase the gift card.

End of Block: R2_offer

Start of Block: R2_choice

Round 2

Please choose **one** of the following options as your bonus payment:

◯ A $10 cash payment

◯ A $10 Walmart gift card and the remaining ${10-price2} in cash (the change)

Note that you will not be able to change this decision once you proceed to the next screen.

End of Block: R2_choice

Start of Block: R2_erase (displayed for condition E)

Round 2

Before proceeding to the next round, please indicate whether you would like the algorithm to use the information about your decision to set prices in future rounds.

◯ use information

◯ do not use information

End of Block: R2_erase

Start of Block: R3_offer

Round 3

$10 has been added to your bonus payment for this round. Based on the information we have about you, in this round the algorithm is offering you the $10 Walmart gift card **for a price of ${price3}**.

Display This Question:

If treatment = cookies

Before making your choice, please indicate whether you would like the algorithm to use the information about your decision to set prices in future rounds.

○ use information

○ do not use information

Please proceed to the next screen to decide whether you want to use some of your bonus payment to purchase the gift card.

End of Block: R3_offer

Start of Block: R3_choice

Round 3

Please choose **one** of the following options as your bonus payment:

○ A $10 cash payment

○ A $10 Walmart gift card and the remaining ${10-price3} in cash (the change)

 

Note that you will not be able to change this decision once you proceed to the next screen.

End of Block: R3_choice

Start of Block: R3_erase (displayed for condition E)

<u>Round 3</u>

Before proceeding to the next round, please indicate whether you would like the algorithm to use the information about your decision to set prices in future rounds.

- ◯ use information
- ◯ do not use information

End of Block: R3_erase

Start of Block: notice (displayed before rounds 1, 2 or 3, for conditions 1R, 2R and R3, respectively)

Notice: you have been chosen to play {rounds} rounds. This means that the next round will be your last one!

End of Block: notice

Start of Block: R4_offer

<u>Round 4</u>

$10 has been added to your bonus payment for this round. Based on the information we have about you, in this round the algorithm is offering you the $10 Walmart gift card **for a price of ${price4}**.

Display This Question:

If treatment = cookies

Before making your choice, please indicate whether you would like the algorithm to collect the information about your decision in this round.

○ use information

○ do not use information

Please proceed to the next screen to decide whether you want to use some of your bonus payment to purchase the gift card.

End of Block: R4_offer

Start of Block: R4_choice

Round 4

Please choose **one** of the following options as your bonus payment for this final round in which the algorithm sets the price:

○ A $10 cash payment

○ A $10 Walmart gift card and the remaining ${10-price4} in cash (the change)

 

Note that you will not be able to change this decision once you proceed to the next screen.

End of Block: R4_choice

Start of Block: R4_erase (displayed for condition E)

Round 4

Before proceeding, please indicate whether you would like the algorithm to collect the information about your decision in this round.

○ use information

○ do not use information

End of Block: R4_erase

Start of Block: Last_round_feedback

Your choice of bonus payment for this last round has been recorded. Please proceed to the next screen to continue to the final stage of the study.

End of Block: Last_round_feedback

Start of Block: R5_instructions

Final Stage

In this final stage of the study, you will have another chance to purchase a $10 Walmart gift card, but the algorithm will no longer set the price for it. Instead, we will ask you to tell us the maximum amount you are willing to pay for the gift card. In other words, we would like to know how much the $10 Walmart gift card is worth to you. To ensure you provide the most accurate estimate, we will auction the gift card to you in a manner that will make it beneficial for you to report the true amount that the gift card is worth to you.

As in previous rounds, you are first given $10. You will then be asked to state the highest amount you are willing to pay for the gift card. However, the price you will pay for the gift card **will not be** the amount you report. Instead, the computer will randomly draw a number between 0 to 10. If the drawn number is equal to or lower than the number you reported, you will receive the gift card for a price equal to the drawn number (**not** the amount you reported), and the remaining sum (the change) will be added to your bonus in cash. If the drawn number is higher than the amount you reported, the gift card will not be purchased, and the bonus will be $10 in cash.

Notice that you cannot secure a better outcome for yourself than by reporting the true value of the gift card to you. If the drawn number is lower than the value of the gift card you report, you benefit from receiving the gift card at a lower price than it is worth to you. For example, if you report that the most you are willing to pay is $7 and draw the number 3, then you will only pay $3 for the gift card. Conversely, if the drawn number is higher than the value of the gift card you report, you will not buy it. For example, if you report that the most you are willing to pay is $7 and draw the number 9, then you will not be able to buy the gift card and the bonus will be $10 in cash.

In the next screen, we will ask you a few questions to ascertain that you have read these instructions and understand the mechanisms of the study.

End of Block: R5_instructions

Start of Block: R5_comprehension

Please answer the following questions, which are meant to ascertain your comprehension of the instructions in the previous screen. Each question has only one correct answer, and you must answer the questions correctly using the drop-down menus to be able to proceed to the next screen. Please feel free to return to the previous screen to re-read the instructions.

What will be your bonus payment if you report that the most you are willing to pay for the gift card is $7 and the computer draws the number 3?

What will be your bonus payment if you report that the most you are willing to pay for the gift card is $8 and the computer draws the number 9?

What will be your bonus payment if you report that the most you are willing to pay for the gift card is $6 and the computer draws the number 6?

What will be your bonus payment if you report that the most you are willing to pay for the gift card is $9 and the computer draws the number 2?

<options for all 4 questions:>

- ○ only $10 in cash
- ○ a $10 Walmart gift card and $0 in cash
- ○ a $10 Walmart gift card and $1 in cash
- ○ a $10 Walmart gift card and $2 in cash
- ○ a $10 Walmart gift card and $3 in cash
- ○ a $10 Walmart gift card and $4 in cash
- ○ a $10 Walmart gift card and $5 in cash
- ○ a $10 Walmart gift card and $6 in cash
- ○ a $10 Walmart gift card and $7 in cash
- ○ a $10 Walmart gift card and $8 in cash
- ○ a $10 Walmart gift card and $9 in cash
- ○ a $10 Walmart gift card and $10 in cash

End of Block: R5_comprehension

Start of Block: R5_choice

Please report the maximum amount (in dollars) that you are willing to pay for a $10 Walmart gift card from the drop-down menu. Remember that this is not the price you will pay for it. In the next screen, the computer will randomly draw a price, and you will only purchase the gift card if that price is lower than the amount reported here.

- ○ $0
- ○ $1
- ○ $2
- ○ $3
- ○ $4
- ○ $5
- ○ $6
- ○ $7
- ○ $8
- ○ $9
- ○ $10

Note that you will not be able to change this decision once you proceed to the next screen.

End of Block: R5_choice

Start of Block: R5_result

The computer has randomly drawn the number {Random_Num}.

Display This Question:
If Random_Num <= {WTP}

Because ${Random_Num} is equal to or lower than your reported amount of ${WTP}, the gift card has been purchased for you and your bonus for this round is the $10 Walmart gift card and the remaining ${10-Random_Num} in cash

Display This Question:

If Random_Num > {WTP}

Because ${Random_Num} is higher than your reported amount of ${WTP}, you cannot purchase the gift card and your bonus for this round is $10 in cash.

End of Block: R5_result

Start of Block: demographics

Before concluding, please provide us with the following information about yourself. Providing this information is required for completing this study, but it also helps us in analyzing your response with the goal of improving the regulation of consumer markets. As mentioned above, we are a group researchers, and we do not collect any information that could identify you by anyone (including ourselves).

In a few words, why did you choose to either purchase or not purchase the gift card in the different rounds? Was there any reason other than the price in that round?

____

In a few words, what is your best guess regarding the goal of this study?

____

What is your age?

What best describes your gender?

- ○ Male
- ○ Female
- ○ Transgender Male
- ○ Transgender Female
- ○ Non-Binary

○ Other

What best describes your race and ethnicity?

○ Non-Hispanic White/Caucasian

○ African American

○ Hispanic

○ Asian

○ Native American

○ Native Hawaiian/Pacific Islander

○ Other

In which state do you currently reside?

Were you born in the United States?

○ Yes

○ No

Which of these describes your annual income last year?

○ $0

○ $1 to 9,999$

○ $10,000 to 24,999$

○ $25,000 to 49,999$

- $50,000 to 74,999$
- $75,000 or more

What is the highest level of school you have completed or the highest degree you have received?

- Less than high school degree
- High school graduate (high school diploma or equivalent including GED)
- Some college but no degree
- Associate degree in college (2-year)
- Bachelor’s degree in college (4-year)
- Master’s degree
- Doctoral degree
- Professional degree (JD, MD)

What best describes your occupation?

- Management Occupations
- Business and Financial Operations Occupations
- Computer and Mathematical Occupations
- Architecture and Engineering Occupations
- Life, Physical, and Social Science Occupations
- Community and Social Service Occupations
- Legal Occupations

- Educational Instruction and Library Occupations
- Arts, Design, Entertainment, Sports, and Media Occupations
- Healthcare Practitioners and Technical Occupations
- Healthcare Support Occupations
- Protective Service Occupations
- Food Preparation and Serving Related Occupations
- Building and Grounds Cleaning and Maintenance Occupations
- Personal Care and Service Occupations
- Sales and Related Occupations
- Office and Administrative Support Occupations
- Farming, Fishing, and Forestry Occupations
- Construction and Extraction Occupations
- Installation, Maintenance, and Repair Occupations
- Production Occupations
- Transportation and Material Moving Occupations
- Student
- Other

---

how many hours per week do you spend on online shopping websites on average?

- None
- Some but less than one (1) hour

○ One (1) to three (3) hours

○ More than three (3) hours

To the best of your estimate, what portion of your shopping expenditure (apparel, groceries, furniture, electronics, etc.) is spent online?

○ None (0%)

○ 1% to 20%

○ 21% to 40%

○ 41% to 60%

○ 61% to 80%

○ 81% to 100%

What is your best estimate regarding the number of times you shop at Walmart in an average month?

○ 0 times

○ 1-2 times

○ 3-4 times

○ 5 times or more

Please make sure to answer all questions before submitting your responses.

End of Block: demographics