Cover letter: REStat fellowship

Jose Tudon*

May 1, 2025

Date of PhD: 2018.

Personal website: jtudon.com

My research spans environmental and health economics, with a focus on applying structural and empirical tools to understudied questions in emerging markets and digital economies. As a researcher based in Mexico, I bring both institutional and geographic diversity to the NBER community. My work seeks to broaden the scope of applied microeconomics by introducing new datasets, overlooked populations, and novel empirical strategies that address urgent global policy challenges.

One line of research develops a scalable, economics-based methodology to quantify the full environmental footprint of cryptocurrency mining. While most existing studies focus narrowly on Bitcoin and rely on hardware-specific assumptions, I estimate the energy consumption and greenhouse gas emissions of over 7,000 active cryptocurrencies. The results reveal that, although Bitcoin dominates, smaller coins collectively account for substantial emissions, challenging the idea that environmental harms are concentrated in a few major platforms. This work contributes to

*Department of Business Administration, ITAM. Email: jtudon@itam.mx.

energy and environmental economics by offering a comprehensive framework for monitoring sustainability in blockchain systems and informing future regulation.

A second project, situated in the field of health and public economics, explores healthcare consumption under uncertainty. Using rich administrative data from a private cataract surgery provider in Mexico serving low-income patients, we estimate a dynamic model in which individuals are forward-looking and learn from their first surgical experience before deciding on a second. Counterfactual simulations show that price adjustments—rather than persuasive marketing—are more effective in increasing uptake and welfare. This work provides practical guidance on how healthcare providers and policymakers can design budget-neutral pricing strategies to expand access among poor or underinsured populations.

What I hope to contribute to the NBER Summer Institute is a perspective grounded in research on low- and middle-income contexts, guided by questions of access, sustainability, and behavioral responses to uncertainty. Coming from an institution and region historically underrepresented in the NBER community, I offer first-hand insight into the economic challenges faced by populations often absent from mainstream datasets.

At the same time, I hope to benefit from exposure to the latest empirical and theoretical developments in applied microeconomics, deepen collaborations across fields, and gain feedback that will sharpen the policy relevance and methodological rigor of my work.

I am especially interested in participating in the Energy and Environmental Economics and Health Economics sessions, where my ongoing projects are most closely aligned and where I see the greatest opportunity for intellectual exchange and growth.

JOSÉ TUDÓN

jtudon@itam.mx
jtudon.com

Office Contact Information

Department of Business Administration,
ITAM
Av. Camino a Santa Teresa 930
Mexico City 10700, Mexico
(+52) 55-5628-4000

Education

Ph.D., The University of Chicago, 2012 to 2018
Thesis Title: "Congestion v Content provision in Net Neutrality: The case of Amazon's Twitch.tv"

M.A., Economics, The University of Chicago, 2014
B.A., Economics, ITAM, Mexico City, double major with Applied Mathematics, 2011

References:

Ali Hortaçsu	*Michael Dinerstein*	*Brent Hickman*
The University of Chicago	The University of Chicago	Washington University in St. Louis
773-702-5841	773-795-2937	613-533-2250
hortacsu@uchicago.edu	mdinerstein@uchicago.edu	hickmanbr@gmail.com

Research and Teaching Fields

Primary field: Industrial Organization.
Secondary fields: Applied microeconomics, applied theory.

Academic Employment

Assistant Professor, ITAM Business School, Fall 2018 - present.

Publications

1. "Prioritization vs. congestion on platforms: evidence from Amazon's Twitch.tv." in *The RAND Journal of Economics*, Volume 53, 328-355, 2022.
2. "Distilling network effects from Steam" in *Quantitative Marketing and Economics*, Volume 20, 293–312, 2022.
3. "Can Price Dispersion be supported solely by Information Frictions?" in *Economic Theory Bulletin*, Volume 9, 75-90, 2021.

4. "Perception, Utility, and Evolution," in *Economic Theory Bulletin*, Volume 7, 191–208, 2019.
5. "Pay-what-you-want because I do not know how much to charge you" in *Economics Letters*, Volume 137, 41-44, 2015.
6. "The profitability and performance of U.S. regional banks using the predictive focus of the 'fundamental analysis research' " with Glenn Growe, Marinus DeBruine and John Y. Lee in *Advances in Management Accounting*, Volume 24, 191-239, 2014.
7. "A simple search model to avoid Diamond paradox" (in Spanish) in *Gaceta de Economía* núm. 29, año 17, ITAM, 2011.
8. "An escape to Diamond paradox" with José García (in Spanish) in *Gaceta de Economía* núm. 28, año 16, ITAM, 2010.

Working Papers

1. "Eye-opening products: Uncertainty and surprise in cataract surgery outcomes" with Emilio Gutierrez and Adrian Rubli.
2. "Quantifying the energy and carbon costs of all cryptocurrencies: A scalable economic approach."
3. "Promotions measured with error: Implications for health policy and firm behavior" with Emilio Gutierrez and Adrian Rubli.
4. "A Memo on the Proof-of-Stake Mechanism" with George Gui and Ali Hortaçsu.

Honors, Scholarships and Fellowships

Víctor L. Urquidi Prize in Economics 2022, 2nd place	2023
Conacyt SNI, level 1 member	2023–
Conacyt SNI, level C member	2020–2022
NET Institute Research Grant Recipient	2017
Finalist, Illinois Economics Association's Graduate Paper Competition	2017
Sherwin Rosen Fellowship, The University of Chicago	2017
Presidential Scholarship, Division of Social Sciences, The University of Chicago	2012–2017
Outstanding Author Contribution, Emerald Literati Network Awards for Excellence	2015
Econometric Society World Congress 2015 Travel Grant Award	2015
Stony Brook Center for Game Theory Travel Award	2015
Graduate Student Affairs Travel Award, The University of Chicago	2015

Teaching Experience

Courses at ITAM:	
Data-driven decision-making for businesses (Undergrad, Master)	Spring 2019–
Courses at UAT:	
Data and business decision-making (Ph.D.)	Falls 2019, 2022
Teaching Assistant at The University of Chicago:	
Introduction to Microeconomics (Undergrad) for John List and Anya Samek	Spring 2015
Evolutionary Game Theory (Ph.D.) for Balázs Szentes	Springs 2015, 2017

Elements of Economic Analysis II (Undergrad) for Richard Van Weelden Winter 2016
Analysis of Collective Decision Making (Undergrad) for
Richard Van Weelden Spring 2016
Time Series Econometrics (Undergrad) for James Marrone Spring 2016
Microeconomics (MBA) for Seth Zimmerman Fall 2016
Introduction to Industrial Organization (Undergrad) for Michael Dinerstein Winter 2017

Conference Presentations

North American Winter Meeting of the Econometric Society, San Antonio, TX Jan 2024
75th EEA - 38th ESEM congress, Barcelona, Spain Aug 2023
50th EARIE Annual Conference, Rome, Italy Aug 2023
36th Jornadas de Economía Industrial, Las Palmas de Gran Canaria, Spain Sept 2022
49th EARIE Annual Conference, Viena, Austria Aug 2022
2022 Econometric Society Australasia Meeting (Hybrid) Jul 2022
2022 Asian Meeting of the Econometric Society, Shenzhen, China (Hybrid) Jun 2022
INFORMS Annual Meeting, Anaheim, CA (Virtual) Oct 2021
48th EARIE Annual Conference, Bergen, Norway (Virtual) Aug 2021
Asian Meeting of the Econometric Society, Miri, Malaysia (Virtual) Jun 2021
19th International Industrial Organization Conference, Virtual Apr 2021
4to Congreso de Inv. en Comportamiento Org. e Indiv., Tampico, Mexico Oct 2019
46th EARIE Annual Conference, Barcelona, Spain Aug 2019
16th Annual International Industrial Organization Conference,
Indianapolis, IN Apr 2018
47th Annual Meeting, Illinois Economics Association, Chicago, IL Oct 2017
2nd Annual Digital Information Policy Scholars Conference,
George Mason U. Apr 2017
11th World Congress of the Econometric Society, Montréal, Canada Aug 2015
26th International Conference on Game Theory, Stony Brook University Jul 2015
13th Annual International Industrial Organization Conference, Boston, MA Apr 2015

Seminar Presentations

Olin Business School, WashU 2022
Rice University 2020
US Federal Communications Commission 2018
University of Mannheim, Cornell University, Clemson University 2018
CIDE, Banco de México, ITAM 2018
Applied Microeconomics Workshop, The University of Chicago 2017
Industrial Organization Workshop, The University of Chicago 2014, 2016, 2017
Economic Theory Workshop, The University of Chicago 2014, 2015, 2016

Non-academic presentations

"*Rich in knowledge (Rig på viden)*" podcast, episode 96: Pay-what-you-want pricing, in `spotify.com`, December 2022.
"*What to expect in 2024: How to make the most from your media strategy*," webinar organized by Nielsen IBOPE Mexico, February 2024.

Professional Service

Reviewer for: *Austrian Science Fund (FWF), Journal of Political Economy, Journal of Public Economics, Journal of Behavioral and Experimental Economics, Latin American Economic Review, Journal of the Knowledge Economy.*
Editor-in-Chief of the *Gaceta de Economía* January 2011– June 2011
Editor of the *Gaceta de Economía* 2009–2011

Research Assistance Experience

Research Assistant to Ali Hortaçsu, The University of Chicago Fall 2013– Spring 2018
Research Assistant to Enrique Seira, ITAM January 2011– June 2012
Research Assistant to Miguel Iraola, ITAM January 2011– June 2011

Students

Pre-doc advising and doctoral placement:
David Ontaneda (PhD Econ, U Michigan)
Adrián Martínez (PhD Econ, UCSD)
Jesús López (PhD Econ, CUNY)

Languages

English (fluent), Spanish (native), Italian (modico), Chilean spanish (basic).

Personal Information

Born June 13th, 1987 in Ciudad Madero, Tamaulipas, México. He/Him. Married. Father of three daughters, born in 2018, 2020, and 2023. Mexican citizen.

Eye-opening products: Uncertainty and surprise in cataract surgery outcomes*

Emilio Gutierrez† Adrian Rubli‡ José Tudón§

March 7, 2025

Abstract

When consuming experience goods, individuals face uncertainty about the benefits, which are revealed only through consumption. This uncertainty may lead to under-consumption, especially when repeat interactions are limited. Cataracts, which typically develop in both eyes among older adults, can only be treated through surgery, though uptake is generally low. We develop a dynamic demand model in which consumers are forward-looking and learn from their first consumption before deciding on a second interaction. Using data from a large private cataract surgery provider in Mexico that primarily serves low-income patients, we estimate demand and identify structural parameters—including price elasticities and the value of the uncertain shock—by leveraging variation in sales agent interactions. We then simulate counterfactual scenarios that attempt to balance the firm's revenue objectives with its mission to expand healthcare access. These simulations suggest that budget-neutral price adjustments are more effective than persuasive advertising at increasing both sales and consumer welfare, highlighting the potential for pricing policies even when repeat interactions are limited.

*We thank Michael Dinerstein, Brent Hickman, Felipe Brugués, Francisco Garrido, and conference and seminar participants for helpful comments. We acknowledge support from the Asociación Mexicana de Cultura. The views and conclusions presented in this paper are exclusively the responsibility of the authors and do not necessarily reflect those of our partner in the field. We are extremely grateful to the staff members at our partner firm for their invaluable help in accessing the data and explaining the institutional details. Felipe García-Moreno and David Sosa provided excellent research assistance. All errors are our own.

†Department of Economics, ITAM. Río Hondo 1, CDMX 01080, Mexico. Email: emilio.gutierrez@itam.mx.

‡Department of Business Administration, ITAM. Email: adrian.rubli@itam.mx.

§Department of Business Administration, ITAM. Corresponding author. Email: jtudon@itam.mx.

Keywords: demand estimation; uncertainty; experience goods; treatment choices; cataract surgery.

1 Introduction

In markets for experience goods, consumers are ex-ante unsure about product characteristics, such as quality (Nelson, 1970). This ex-ante uncertainty might be especially salient in markets where individuals or households only tend to consume one of these products at a time, and often consider purchases to be long-lasting, as is the case with durable goods, like electric vehicles or housing. However, many experience goods allow for potentially repeated interactions, such as prescription drugs, education services, entertainment, or some elective surgeries. In these cases, the consumer's initial take-up decision may account not only for inherent uncertainty but also for the fact that information about idiosyncratic benefits is revealed before making subsequent choices. As a result, take-up may increase or decrease depending on the level of uncertainty, risk preferences, and the option value of the first choice. These effects may also be influenced by the number of potential repeated interactions, which, in some settings, are limited and known in advance.

This paper focuses on the market for cataract surgeries, a particularly underused yet important experience good in the healthcare industry, because forgoing treatment implies a lower quality of life or, even potentially, worse health outcomes (Keel et al., 2021; Ehrlich et al., 2021). We estimate a structural model of demand for cataract surgeries, exploiting patient-level data from a large private provider in Mexico City. We explicitly incorporate the fact that undertaking surgery for the first eye reveals idiosyncratic information about the benefits from surgery, allowing forward-looking patients to have more information before having to decide on getting surgery on their second eye. We recover estimates of patients' price elasticity for each surgery (first vs second) as well as individual-specific uncertainty parameters. We then use our estimates to evaluate counterfactuals that consider policies proposed in the medical literature as a means for increasing the number of cataract surgeries.

Cataracts occur when the eye's natural (and normally clear) lens becomes cloudy, due to a breakdown of its proteins. Cataracts lead to vision problems, ranging from blurry eyesight to complete loss of sight. Age is the leading risk factor, although co-morbidities and risky behaviors (such as smoking) may also

increase their likelihood (Miller et al., 2022). The majority of patients develop cataracts to some degree in both eyes. Surgery to replace the clouded lens with an artificial intra-ocular lens is the only available treatment (Miller et al., 2022). Physicians tend to recommend surgery on both eyes, but perform the surgeries in a sequential manner to minimize complications and inconveniences during post-operative care (Henderson and Schneider, 2012).

While cataracts are a common condition around the world, patients remain massively undertreated—particularly, in low- and middle-income countries (Lansingh et al., 2010; Congdon and Thomas, 2014). For instance, in Mexico, about 350,000 new cases are diagnosed each year, of which only about 50% undergo surgery, with similar numbers across the developing world.[1] Some commonly identified barriers to surgery include prices, access, and uncertainty (Lewallen and Courtright, 2000; Syed et al., 2013).

Cataract surgeries are an ideal setting for studying how uncertainty and revealed information affect initial and consequent product take-up of experience goods with (limited) repeated interactions. First, because most patients develop cataracts in both eyes, the number of potential repeated interactions is ex-ante known to the patient.[2] Second, it is very uncommon for patients to undergo surgery in both eyes simultaneously (in our data, none do so). Lastly, in our setting, patients are not given the option of scheduling surgery for both eyes simultaneously; instead, they must schedule and pay for each surgery sequentially, guaranteeing that the first surgery's benefits are realized before the patient must decide on the second surgery. We use the term "learning" to refer to this unknown information that is revealed after consumption. Taken together, these features allow us to effectively model this decision-making process in two stages, with information revealed before having to decide on the second surgery (conditional on having chosen to operate in the first stage).

We analyze the decision-making process of patients at a low-cost private provider in Mexico City specializing in cataract surgery. We obtain patient-level records that allow us to observe cataract diagnoses, subsequent price quotes, and whether the patient purchases a surgery. Our data contain all patients whose first contact with this provider occurred during 2018, and whom we observe

[1] See `excelsior.com.mx` and Cataract surgical rates (2017). In this paper, urls are truncated, but their hyperlinks are not. Urls last accessed June, 2023.

[2] This is not always the case for other types of healthcare goods. For example, patients with vascular disease may have varying numbers of affected arteries at varying degrees of deterioration, which may or may not worsen with age. Therefore, patients might need anywhere from one to many angioplasties in a relatively short time span to manage this condition. See, for example, Gheini et al. (2021).

over multiple visits to the clinics during 2018 and 2019. Aside from the features outlined above, we leverage the fact that prices are not homogeneous across patients. After the diagnosis, patients are assigned to a sales agent based on availability, who then proposes a price quote from a menu of discount options.

We begin by exploring the potential for learning in cataract surgeries through an out-of-sample survey, an empirical analysis of our data, and a test based on a simple model. Building on these insights, we then develop a model where a patient decides sequentially whether to undergo surgery for each eye, conditional on her current information set. From her point of view, there is an uncertain component in the outcome of the first operation, which is only revealed after experiencing the first surgery. Then, all (knowable) information is known to the patient before having to decide on the second surgery. This setup implies consumers have an option value from the first surgery. In our estimation, we deal with endogenous prices with a control function that uses sales targets as an instrument, and we deal with potential selection into returning after the first surgery by simulating announced prices for the second surgery for patients who did not come back for a quote. Throughout, we allow for decreasing marginal utilities for the second operation, flexible risk preferences, and control for patients' income, health, and other demographic characteristics. Thus, we identify the magnitude of the uncertainty shocks from discrepancies in estimated coefficients between the first and second surgery.[3]

We find elastic demand curves, and our estimates show that demand elasticities for the first operation are larger in absolute value than those for the second surgery. We also obtain a heterogeneous distribution of estimated uncertainty in the population, which suggests that the option value of the second surgery has an important role in the decision-making process of patients. We also back out a money metric of consumer surplus indicating quite a bit of variation across patients.

Given our results, we explore potential strategies that firms or providers could implement to increase take-up and patient welfare. Using our estimated parameters, we simulate counterfactual marketing and pricing strategies to assess their impact on cataract surgery uptake. First, we consider interventions

[3]Randomized trials have found that the second surgery leads to significant improvements in both visual acuity and quality of life (Javitt et al., 1995; Laidlaw et al., 1998; Castells et al., 2006). However, there is little evidence on whether the marginal utility decreases or not. One observational study found an only slightly lower marginal utility of the second surgery relative to the first (Busbee et al., 2003). However, another observational study found that marginal utility was increasing among US veterans (Shekhawat et al., 2017).

related to the uncertainty parameter in the form of information provision (or persuasive advertising). Experimental studies that attempt to fully eliminate uncertainty have found mixed results (Liu et al., 2012), while awareness campaigns in which a "champion" (i.e., someone with a positive result) informs about potential outcomes have been more favorable (Mailu et al., 2020). In our exercise, we consider that the champion reveals a particular value for the uncertain shock—which the patient takes as true—in an attempt to persuade patients about potential outcomes. Our simulations show that this intervention might be welfare-improving, as long as the size of the revealed information shock is large enough. That is, the champion must reveal a credible, sizable shock. If the firm is also covering the cost of the campaign, this strategy may not be financially viable.

Our second set of counterfactual exercises examines revenue-neutral price adjustments, where the price of the first surgery is subsidized while the second is taxed, keeping the firm's total revenue unchanged. While the firm may aim to maximize profits, this exercise serves as a benchmark for the range of pricing strategies it may consider. A priori, it is not obvious if total surgeries would go up or down, because it depends on relative price elasticities. Across a range of symmetric and asymmetric price changes, we consistently find large welfare gains: both consumer surplus and overall take-up increases. To put a specific number: a 10% revenue-neutral price change increases demand by about 10%. That is, this simple economic policy generates 10% more cataract surgeries in the market.

Overall, these exercises suggest that persuasive advertising that reduces uncertainty will not be as effective, unless the firm is able to truly convince potential patients that their outcome will be very positive. Instead, by explicitly considering the dynamic link and the option value due to uncertainty that is revealed, we show how—at least in this setting—welfare-improving price changes can be implemented due to the size and heterogeneity of the uncertainty. Although the number of potential interactions is fixed in this context, this intuition extends to other experience goods where the number of interactions is uncertain (e.g., angioplasties for vascular disease) or where the duration of consumption is ex-ante unknown (e.g., antidepressants).

1.1 Related literature

Our paper speaks to various strands of literature. First, we build on a vast literature of consumer learning. In the canonical consumer learning model, consumers face uncertainty about product attributes, hold prior beliefs, and update them through consumption. These models typically assume that consumption provides a noisy signal—due to inherent product variability or consumer-specific match quality—which may become clearer with repeated interactions. The key intuition is that trying something new is risky, but it also generates valuable information. Consumers thus weigh the potential benefits of learning against the risks of uncertainty. See Ching, Erdem and Keane (2017) for a review. In our model, the number of potential interactions is fixed at two. Moreover, without loss of generality, we assume that the first surgery reveals all relevant information, or at least all information the consumer might use when deciding whether to undergo the second and final surgery.

We add to a long-standing literature in industrial organization analyzing dynamics in experience goods markets (Bergemann and Välimäki, 2006; Gowrisankaran and Rysman, 2012; Jing, 2011; Yu, Debo and Kapuscinski, 2016). However, unlike many of these settings, ours is one with a limited and small number of repeated interactions, which may affect the capacity of the firm to adapt and hinder customer reactions to these dynamics. This feature may be relevant in other settings as well, such as durable goods markets.

Related work on the role that uncertainty plays in demand has also focused on how providing additional external information—for instance, in the form of expert advise or customer reviews—might affect product demand. Studies in this area have analyzed, among others, negative book reviews (Berger, Sorensen and Rasmussen, 2010), movie critics (Reinstein and Snyder, 2005), and expert opinion labels for wine (Hilger, Rafert and Villas-Boas, 2011). Moreover, a related literature has further explored the effects of free trials before purchasing on consumption decisions (Foubert and Gijsbrechts, 2016; Sunada, 2020).

Our paper is also related to the health economics literature attempting to understand dynamic treatment choices under uncertainty (Dardanoni and Wagstaff, 1990; Manski, 2018). In particular, it has been shown that in low- and middle-income countries, demand for pharmaceutical treatments is inelastic while demand for diagnoses is more elastic (Dupas and Miguel, 2017). Our results echo this finding: once patients are aware of the benefits, they respond more inelastically. Our findings on the importance of the learning effect is also consistent

with evidence on the adoption of health products in these developing country settings (Dupas, 2014; Oster and Thornton, 2012).

Other work in this area has focused on search and learning costs for pharmaceutical products, for instance, in the context of generic prescription drugs (Ching, 2010), anti-ulcer drugs (Crawford and Shum, 2005), antidepressants (Dickstein, 2021), and flu shots (Maurer and Harris, 2016). Some studies have focused on the role of advertising to reduce uncertainty on both patient and physician demand for drugs (Narayanan and Manchanda, 2009; Ching and Ishihara, 2010; Ching et al., 2016; Wosinska, n.d.). Finally, studies in the medical literature have documented the uncertainty that patients face before surgery and the potential consequences for treatment choice, recovery, and prognosis for these patients (Lien et al., 2009; Zhang, 2017; Lee et al., 2020; Çakir, Küçükakça Çelik and Çirpan, 2021). Our paper adds to this broad literature by identifying the option value of the first round of consumption in a context where the number of repeated interactions is fixed. Indeed, in a setup with only two interactions, it is not obvious that subsidizing initial take-up is an efficient use of resources, although here, with this level of uncertainty, it actually improves welfare.

Lastly, we contribute to the (mostly) medical literature exploring why take-up rates of cataract surgeries are low. Although cataracts are an important cause of blindness in advanced age worldwide, many patients do not undergo surgery. Experimental studies have shown that prices are an important barrier (Zhang et al., 2013), but there is little consensus on the impact of other factors such as information, uncertainty, and peer effects, among others (Mailu et al., 2020; Adhvaryu et al., 2020). Our paper innovates on these experiments by focusing on the dynamic problem inherent to cataract surgeries and the fact that previously unknown information is revealed after the first surgery. Furthermore, our counterfactual exercises show important ways in which take-up may be increased across settings.

2 Context

Cataracts are a condition where the lens of the eye becomes clouded, leading to important declines in eyesight. Age is the biggest risk factor, in particular, due to the cumulative effects of ultraviolet radiation or oxidative damage (Hashemi et al., 2020). Other important risk factors include obesity, high blood pressure, and diabetes. Because cataracts are "an inevitable side effect of aging" (Hashemi

et al., 2020), virtually all patients develop cataracts to some degree in both eyes. According to the National Eye Institute, around 45% of Americans ages 75-79 are affected, as well as over 60% of those ages 80 and over.[4]

While early symptoms may improve with glasses, advanced cataracts require surgery to replace the lens with an artificial one. Most surgeries nowadays use a technique called phacoemulsification, whereby the eye's internal lens is emulsified and vacuumed out of the eye. Alternatively, the doctor may make a series of small incisions to remove the lens; usually, the small incision surgery is more appropriate for worse cataracts. An artificial lens, made of various materials, is then placed in the eye. Generally, the ophthalmologist decides the type of surgery and lens based on the medical and physiological needs of the patient. Importantly, physicians recommend surgery on both eyes to avoid binocular inhibition or discomfort from having dissimilar vision in the two eyes (Miller et al., 2022).[5]

Access to these surgeries in Mexico is hindered by characteristics of the healthcare system, which is a mix between private and public providers (OECD, 2016). The government supplies healthcare coverage to individuals through its own network of providers. This public system is mostly free of charge, but is plagued by long waiting times and heterogenous quality. Alternatively, patients may visit private providers, though low private insurance rates imply that most of these services are paid for out-of-pocket. This results in relatively high price elasticities of demand. Hence, large segments of the population do not have access to private care given their high prices and the lack of health insurance.

Estimates suggest that 30-40% of individuals in Mexico have cataracts, with 350,000 new cases per year. With diabetes cases on the rise, cataract rates in non-elderly populations are increasing as well.[6] Although cataract surgery is covered by the public healthcare system, long waiting times hamper timely access to treatment.[7] Furthermore, clinical guidelines in the public sector only allow cataract surgery once the patient's eyesight is severely deteriorated, at a much higher threshold of vision loss than the standard of care in developed

[4]See `nei.nih.gov`.

[5]After a second surgery patients generally report better stereopsis (binocular sight), visual acuity, and overall functionality compared to first-eye-only patients; see Miller et al. (2022); Castells et al. (2006); Laidlaw and Harrad (1993); Heemraz et al. (2016).

[6]See `excelsior.com.mx`. Recent estimates from public providers suggest that 15 to 20% of young adults are affected by cataracts in Mexico; see `imss.gob.mx`.

[7]According to information from our partner firm, patients in the public system wait ten months on average for cataract surgery after diagnosis.

nations, like the US.[8] In the private market, recommendations for surgery follow international standards, but most surgeries cost between 1,300 and 1,500 USD per eye, which is equivalent to 1.5-1.7 times the median monthly household income in the Mexico City metropolitan area (ENIGH 2018).[9]

Our partner firm is a large private provider of ocular healthcare based in the Mexico City metropolitan area that opened to the public in 2011. The firm provides various eye care services such as regular check-ups, eye exams, lab analyses, surgery, and an optical store. Their clinics are spread out over 20 locations, with a main clinic in downtown Mexico City, where the majority of surgical interventions are carried out. An important part of the firm's business is diagnosing and operating cataracts, with a target population made up of mostly lower-income patients.

Consumer's journey. When a patient arrives at the facilities of our partner firm, an ophthalmologist evaluates the patient's eyes, orders on-site lab analyses, and performs exams, some of which are conducted by in-house optometrists.[10] The physician diagnoses cataracts by assigning each eye a cataract score ranging from zero (no cataracts) to six (severe cataracts). Surgery is generally recommended for patients with a score of three or higher, signifying blurry eyesight. However, ocular comorbidities and physician practices may also play a role in determining the need for surgery. Once a patient has been diagnosed by the physician and she has deemed that surgery is the recommended treatment, the physician emits a prescription with a recommendation for the type of surgery and type of intraocular lens. This recommendation is based on medical and physiological determinants, such as severity of cataracts, ocular health, and so on. Ophthalmologists do not have discretion over prices and cannot give price quotes. To get a quote, patients are then referred to a sales agent at the clinic.

[8] According to clinical guidelines from IMSS, the largest public provider in Mexico, surgery is required once a patient "has difficulty performing daily activities such as recognizing familiar faces, has reduced mobility, and/or is unable to work and live independently". These guidelines also recognize that this standard is very different from others, such as the UK's NHS (which considers surgery once the patient's vision is blurry or opaque), and that the private market in Mexico may fill this void for patients whose eyesight is not as deteriorated, conditional on their payment capabilities. Note that no single test can objectively define adequate thresholds for cataract surgery, as many considerations are patient-specific and self-reported (Miller et al., 2022). See `imss.gob.mx`.

[9] These quotes are based on posted prices on websites of the most common eye care clinics in Mexico City. Since surgeries are paid mostly out-of-pocket, we were unable to find any systematic statistics on the price of cataract surgeries. However, conversations with our partner firm suggest that these estimates are correct.

[10] For clarity, throughout the paper we use the term "physician" or "doctor" to refer to the ophthalmologist. The diagnosing physician may or may not coincide with the doctor performing the surgery.

Sales agents are assigned based on availability, although it is likely that a returning patient is assigned the same agent as before. The sales agent then generates a price quote for the patient based on various factors, such as patient and surgery characteristics, available discounts, and a certain degree of discretion. Sales agents may offer discounts based on a menu of options that changes over time. However, the agents' commission is based on the sale price, creating an incentive to avoid using the discounts if possible. Patients take the physician's prescription as given, and rarely ask for alternatives; in our data, virtually no patient receives a price quote for a non-prescribed surgery, which is consistent with a literature that finds patients adhere to the physicians' recommendations, specially when patients lack expertise (Finkelstein et al., 2021; Johnson and Rehavi, 2016; Gruber and Owings, 1996).

If the patient chooses to go forward with the surgery, payment plans are discussed and a date is set. Importantly, surgeries are sold as a single-eye product, requiring patients to schedule each surgery independently. Price quotes also correspond only to surgery on one eye, and patients only schedule and pay for surgeries consecutively. All patients in the data consider first a surgery on the eye with worse cataracts, because, in particular, physicians strongly recommend so.[11] For the vast majority of these patients, the outside option is no surgery at all, which we confirmed with a phone survey of a random subsample of patients: about 17% of patients who did not get a surgery with our partner firm ended up getting a surgery elsewhere.

After the surgery, the patient returns, at the physician's discretion, for follow up appointments. If the patient wants surgery for the second eye, the patient is again referred to a sales agent, who gives another price quote for the prescription of the second eye's surgery, and another date is set. On average, patients who undergo both surgeries do so within 74 days of each other.

3 Data

We obtained anonymized patient-level data directly from our partner firm spanning all patient visits from 2018 and 2019.[12] We restrict our attention to new patients in 2018, allowing us to observe repeated interactions with the firm over

[11] Patients may be required to make a small down payment in order to schedule the surgery. The firm may also provide interest-free credit by allowing patients to pay over various installments, although the full cost must be covered by the day of the surgery.

[12] The data and replication files can be accessed here or at the authors' websites.

a span of at least one year. The data contain some observable time-invariant patient characteristics that include gender, age, whether they are covered by private insurance, whether they have access to public healthcare, and their zip code.

We also observe details from all patient visits. For each one, we observe the service provided by our partner firm, any diagnoses made, and all price quotes generated by sales agents. Physician and sales agent identifiers are included for each interaction. This effectively allows us to observe, for each patient visit, all relevant interactions with medical and non-medical staff, and which products were offered to them, at what price, and whether a purchase was made. Observations are therefore at the patient-visit-product level, regardless of whether the product was actually bought.

We focus our attention on patient-visit-product observations related to cataract diagnoses and surgery products with non-missing or duplicated information. We exclude a small number of patients that had three cataract surgeries over this period, a few cataract surgeries that were not cataloged as either phacoemulsification or small incision surgery, and some pro-bono surgeries for which patients were not billed. Lastly, we exclude patients in the top one percent of the age distribution (i.e., aged 91 and older). Overall, we are left with a sample of 3,822 patients and a total of 4,554 patient-quote observations.

In our data, cataracts in each eye are measured by ophthalmologists on a zero to six scale, where zero denotes no cataracts and six denotes the most severe cases. If cataracts were not reevaluated on a particular visit, we assign the cataract score from the patient's previous visit. Although we focus on patients with cataract diagnoses and surgery products, all patients are evaluated and receive a cataract score at least once during this period. As noted above, virtually all patients develop cataracts to some degree in both eyes, because age is the main risk factor (Hashemi et al., 2020). In our data, we find a high correlation in cataract scores between eyes; a one unit increase in the cataract score of the most afflicted eye is associated with an increase of .4 in the healthier eye.

Table 1 shows summary statistics of individuals' characteristics and ocular health measures for patients in our sample and those who visited the clinic for non-cataract related reasons, and are thus not included in our sample. Appendix Table 5 describes these health measures in further detail. Note that 65% of our patient sample ends up having at least one cataract surgery in this time span. Cataract patients are older, less likely to be privately insured, more likely to have access to the public healthcare system, and more likely to not be covered

by any type of healthcare (public or private). As expected, our patients have significantly worse cataract scores—taking the average over all measurements or simply the highest value in this time period. We also observe nine different eyesight measures for each eye which we take as a proxy for ocular health. In general, we find worse measures among the patients in our sample, perhaps due to a correlation of these measures with cataracts or simply the fact that they are older. From a medical perspective, these ocular comorbidities make surgery the only recommended treatment (Lundström et al., 2015). In our estimations, we control for ocular health to allow for different marginal valuations of surgeries, depending on these characteristics.

We observe a total of 4,554 price quotes for cataract surgeries, of which 3,858 correspond to first surgeries and 696 to second surgeries. Of those who received a first quote, 2,468 underwent the surgery, and of those who returned for a second quote, 617 underwent the surgery.

Importantly, our data allow us to proxy for risk aversion. Indeed, we observe all interactions and visits between patients and firm, so we observe how many visits it takes for a patient to obtain a price quote. At each visit, patients may obtain more information or reassurance about the procedure (i.e., information is weakly increasing in the number of visits between the initial diagnosis and obtaining the price quote). We thus make the reasonable assumption that, patients who are more averse register more visits before obtaining a price quote than patients who are less averse, all else equal. Therefore, as a proxy for risk aversion, we use the number of visits between initial diagnosis and obtaining a price quote, which is 3.9 on average in our sample. We further discuss risk aversion in the context of our model below.

For each surgery price quote, we observe the largely exogenous surgical characteristics, which were determined by medical reasons. At 65% of surgery quotes in our sample, phacoemulsification is more common than small incision surgeries, but is also more expensive. However, patient outcomes and complication rates are similar across both methods (Gogate et al., 2005; Riaz, de Silva and Evans, 2013). We also observe the type of artificial lens that replaces the natural lens and if patients pay for additional services (e.g., lab work) at the time of purchase.[13] Lastly, we observe whether the patient bought the surgery offered by the sales agent in the price quote. In our estimations, we include sales agent fixed effects.

[13]Patients with small incision surgeries are only fitted with one type of lens, while the physician has four options available for phacoemulsification.

TABLE 1: Patient summary statistics

Variable	(1) Cataract patients	(2) Non-cataract patients	(3) Difference
Has a cataract surgery	0.65	0	-
	(0.48)		
Age	69.22	46.92	-22.30
	(12.25)	(20.18)	(0.33)
Female	0.61	0.61	0.00
	(0.49)	(0.49)	(0.01)
Private insurance	0.07	0.15	0.08
	(0.26)	(0.36)	(0.01)
Social security	0.22	0.20	-0.02
	(0.41)	(0.40)	(0.01)
Uninsured	0.72	0.68	-0.04
	(0.45)	(0.47)	(0.01)
Right eye cataract potential	2.68	0.43	-2.25
	(1.61)	(1.03)	(0.02)
Left eye cataract potential	2.64	0.44	-2.20
	(1.61)	(1.04)	(0.02)
Right eye maximum cataract potential	3.06	0.56	-2.50
	(1.58)	(1.18)	(0.02)
Left eye maximum cataract potential	3.02	0.56	-2.46
	(1.59)	(1.20)	(0.02)
Right eye far visual acuity count fingers	0.26	0.08	-0.18
	(0.44)	(0.27)	(0.00)
Right eye far visual acuity hand motions	0.11	0.03	-0.09
	(0.32)	(0.16)	(0.00)
Right eye far visual acuity light perception	0.05	0.02	-0.03
	(0.22)	(0.14)	(0.00)
Right eye far visual acuity no light perception	0.01	0.01	-0.00
	(0.10)	(0.09)	(0.00)
Left eye far visual acuity count fingers	0.23	0.08	-0.15
	(0.42)	(0.27)	(0.00)
Left eye far visual acuity hand motions	0.09	0.03	-0.06
	(0.28)	(0.16)	(0.00)
Left eye far visual acuity light perception	0.05	0.02	-0.03
	(0.22)	(0.14)	(0.00)
Left eye far visual acuity no light perception	0.01	0.01	-0.00
	(0.11)	(0.09)	(0.00)
Right eye ampliopia	0.03	0.01	-0.02
	(0.18)	(0.11)	(0.00)
Right eye anisometropia	0.01	0.00	-0.01
	(0.11)	(0.06)	(0.00)
Right eye astigmatism	0.62	0.57	-0.05
	(0.49)	(0.49)	(0.01)
Right eye myopia	0.39	0.34	-0.04
	(0.49)	(0.47)	(0.01)
Right eye presbyopia	0.34	0.23	-0.11
	(0.47)	(0.42)	(0.01)
Right eye hypermetropia	0.23	0.20	-0.03
	(0.42)	(0.40)	(0.01)
Right eye emmetropia	0.01	0.02	0.01
	(0.08)	(0.13)	(0.00)
Left eye ampliopia	0.03	0.01	-0.02
	(0.17)	(0.11)	(0.00)
Left eye anisometropia	0.01	0.00	-0.01
	(0.10)	(0.06)	(0.00)
Left eye astigmatism	0.61	0.57	-0.04
	(0.49)	(0.50)	(0.01)
Left eye myopia	0.37	0.34	-0.03
	(0.48)	(0.47)	(0.01)
Left eye presbyopia	0.34	0.23	-0.11
	(0.47)	(0.42)	(0.01)
Left eye hypermetropia	0.25	0.20	-0.05
	(0.43)	(0.40)	(0.01)
Left eye emmetropia	0.01	0.02	0.01
	(0.08)	(0.13)	(0.00)
Observations	3,822	39,218	43,040

Notes: Standard deviations in parentheses. Patient characteristics for those within our estimating sample (having at least one cataract-related visit) and those not in our sample. Cataract potential is a score based on a 0-6 classification. The maximum cataract potential is the largest score observed during the study period. All other ocular health measures are binary variables. Last column reports a difference-in-means test with standard errors in parentheses.

TABLE 2: Summary statistics by type of surgical product

Surgical method:	SICS	Phaco.	Diff.
Age	71.01	68.43	-2.58
	(11.01)	(12.67)	(0.38)
Female	0.61	0.61	-0.00
	(0.49)	(0.49)	(0.02)
Private insurance	0.05	0.08	0.03
	(0.22)	(0.27)	(0.01)
Social security	0.26	0.19	-0.06
	(0.44)	(0.40)	(0.01)
Uninsured	0.70	0.73	0.03
	(0.46)	(0.44)	(0.01)
Right eye cataract potential	2.72	2.40	-0.32
	(1.66)	(1.70)	(0.05)
Left eye cataract potential	2.76	2.41	-0.36
	(1.69)	(1.66)	(0.05)
Price (MXN)	8903.17	15365.97	6462.80
	(2858.84)	(5430.41)	(149.42)
Observations	1,501	3,053	4,554

Notes: This table shows summary statistics by type of product offered in each quote. Observations are at the patient-quote level. Phacoemulsification and SICS (small incision cataract surgery) refer to the method used by the surgeon. The third column shows a difference-in-means test. During this period, 1 USD = 19.22 MXN.

Table 2 shows summary statistics at the patient-quote level and distinguishes between phacoemulsification and small incision quotes. As noted above, phacoemulsification is about 70% more expensive than small incision surgery. In our sample, phacoemulsification is also associated with slightly less severe cataract scores, younger patients, and those more likely to be privately insured (which in turn signals a higher socioeconomic status). In our estimations, we control for surgical and patient characteristics, including type of insurance as a proxy for income.

3.1 Evidence of learning

In the context of cataract surgeries, the medical literature has documented that patients update their beliefs after the first surgery (Cheung and Sandramouli, 2005; Henderson and Schneider, 2012). Building on this, we first present out-of-sample evidence on learning using survey data from cataract patients in Mexico. We then highlight patterns in our raw data that may be consistent with learning. Finally, we incorporate these insights into a simplified model, which yields a testable prediction to distinguish learning from full information (or no learning). The next section expands this model for our main estimation, but this simplified version provides useful intuition for the interested reader.

3.1.1 Out-of-sample evidence of learning

This section examines a learning process associated with cataract surgeries by presenting suggestive evidence drawn from a different sample than our main patient data. We document that, before their first surgery, cataract patients typically lack full information on a range of relevant dimensions. After the first surgery, however, some of this information is revealed to them, which may then inform decisions on a potential second surgery. Moreover, we provide evidence showing that patients understand that the initial procedure will reveal some of the previously unknown information, and that they recognize ex-ante that this acquired information may prove useful if they later consider surgery for their other eye.

With the aid of a market research firm, we surveyed 94 individuals diagnosed with cataracts in at least one eye, of whom about 37% had undergone at least one surgery. The sample skewed slightly female (55%), was on average ten years younger than our patient sample, and included a likely higher socioeconomic segment (49% held a college degree). To reduce response bias, we randomized both question and answer sequences and the market research firm offered participants monetary incentives.

The survey questions covered three key areas of concern for cataract patients, which we previously identified with the aid of a focus group: concerns about the possibility of less-than-expected improvement in eyesight following surgery, the burden and costs of post-operative care, and the risk of surgical complications.[14]

Among patients who had already undergone at least one surgery, we observed two main insights: first, these patients initially had significant uncertainty on these three key dimensions; second, after undergoing a first surgery, their uncertainty about a second surgery markedly decreased. We term this reduction in uncertainty "learning."

The left-hand side of Figure 1 presents the findings for individuals with one or more surgeries. In the top panel, patients' concerns about eyesight improvement show a distinct shift: before the first surgery, over two-thirds expressed strong concern (dark gray bars), but strong concerns decreased substantially after the first surgery, shifting instead toward expressing only mild concerns. Concerns about post-operative care (middle panel) and surgical complications (bottom panel) followed the same pattern. Notably, the "not at all" option was rarely chosen, likely due to its suggestion of complete absence of concern.

[14]The original survey questions in Spanish can be found in this link: survey_questions.docx.

Next, in Figure 1, right-hand side, for those without any surgery experience, we asked hypothetical questions that revealed a similar trend: significant initial worry before the first surgery, with an expected reduction in strong concerns afterward. This pattern suggests that, even before the first surgery, patients understand its potential value as an informative step for future decisions.[15]

Finally, direct questions further evidenced this learning effect: 97% of those who had already undergone surgery said they had anticipated it would provide valuable information for a possible second surgery, and 91% reported that their concerns diminished after the first surgery. Among respondents without surgery experience, 95% anticipated that a first surgery would provide useful insights, and 81% expected reduced worry about a second surgery after the first one.

In short, these findings underscore the potential role of information gained from the first surgery, highlighting both the reduction in patient uncertainty and patients' awareness of this learning opportunity. This motivates our modeling approach, in which sequential surgery decisions are framed within a limited-information framework; information is revealed after the first surgery, and patients anticipate this learning process when deciding on surgical uptake.

3.1.2 Suggestive evidence of learning in our raw data

We now provide additional evidence in our patient data suggestive of learning. The out-of-sample survey showed that patients had concerns about limited eyesight improvement, post-operative care burdens, and surgical complications. While we cannot observe post-operative care in our data, and surgical complications are extremely rare (with none reported in our sample), we can approximate eyesight improvement beliefs using differences in cataract severity. Cataract surgery fully restores vision in the operated eye by replacing the cloudy lens with a clear synthetic one. Thus, in these exercises, we assume that patients with more severe cataracts "learn more" or "learn differently" than those with milder cases. We do not assume whether patients with severe cataracts are more likely to be positively or negatively surprised by post-surgery vision improvement. Anecdotal evidence suggests that those with slowly progressing, severe cataracts are often surprised by the extent of their vision improvement (e.g., Monet's case in Gruener (2015)). However, high expectations in patients with severe cataracts may also lead to disappointment. Thus, we simply argue here that cataract

[15]To allow for a meaningful comparison, we maintained question wording as consistently as possible between patient groups.

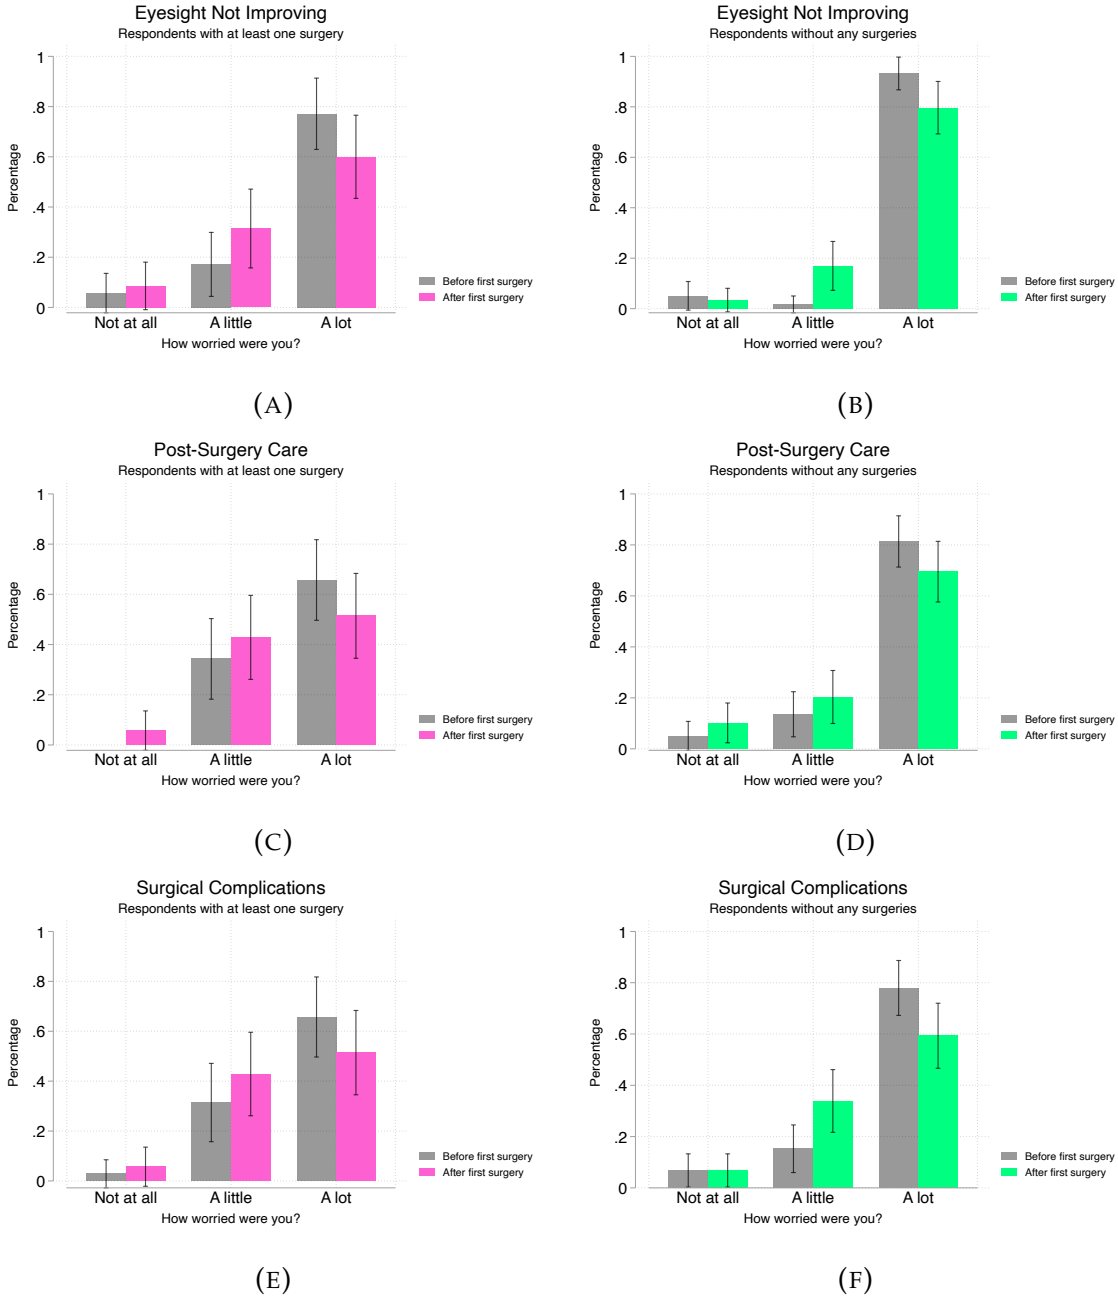


FIGURE 1: Survey results

Note: These figures show results to the online survey of persons with cataract diagnoses. Bars denote the share of respondents for each item. Capped spikes represent 95% confidence intervals. Plots on the left restrict to respondents with at least one cataract surgery ($N = 35$); plots on the right are those who have not yet had any surgery ($N = 59$). Darker gray bars correspond to a question about how respondents felt before their first surgery. Lighter, colored bars are about how respondents felt about a prospective second surgery, after their first one.

severity correlates with learning.

For our first thought experiment in the data, compare patients *before* their first surgery who have zero cataracts in their second eye with patients *after* their first surgery who now have zero cataracts in their first eye. For example, suppose Alice has cataract scores of 4 and 0, while Bob has 6 and 4. After Bob's first surgery, his scores become 0 and 4, making Alice's first surgery comparable to Bob's second. If Alice and Bob are observationally equivalent and there is no learning, their take-up rates should be similar, on average. Note that Alice may still develop cataracts in her second eye over time, meaning she retains some option value from the first surgery.

	Before first surgery		After first surgery
	Alice	Bob	Bob
First eye score:	4	6	0
Second eye score:	0	4	4

Figure 2 compares first-time patients with a zero cataract score in one eye to returning patients, who also have a score of zero in their operated eye. Returning patients are significantly less likely to undergo surgery, and this gap narrows as the cataract score of the remaining bad eye increases. These differences suggest that observationally equivalent patients are not facing the same decision, pointing to learning as a potential explanation. However, factors beyond learning, such as risk aversion, may also play a role (e.g., patients may hesitate to undergo a second surgery if they still perceive significant risks and new information does not alleviate these concerns). To account for this, we next consider a thought experiment that keeps the comparison at the second surgery decision for both Alice and Bob.

Compare two observationally equivalent patients who differ only in their first-eye cataract score. Suppose Alice has a score of 4 and Bob a 6 in their first eye, but both have a 3 in their second eye. Both decide to undergo surgery, such that after their first surgery, their cataract scores are identical (assuming no progression in the second eye). However, Bob, having started with more severe cataracts, may have gained different information from the first surgery than Alice. If no information were revealed through consumption, their take-up rates should be the same, as their marginal utility of second surgery would be identical. Hence, if their likelihood of undergoing the second surgery differs (holding all else equal), it may be attributed to differences in learning.

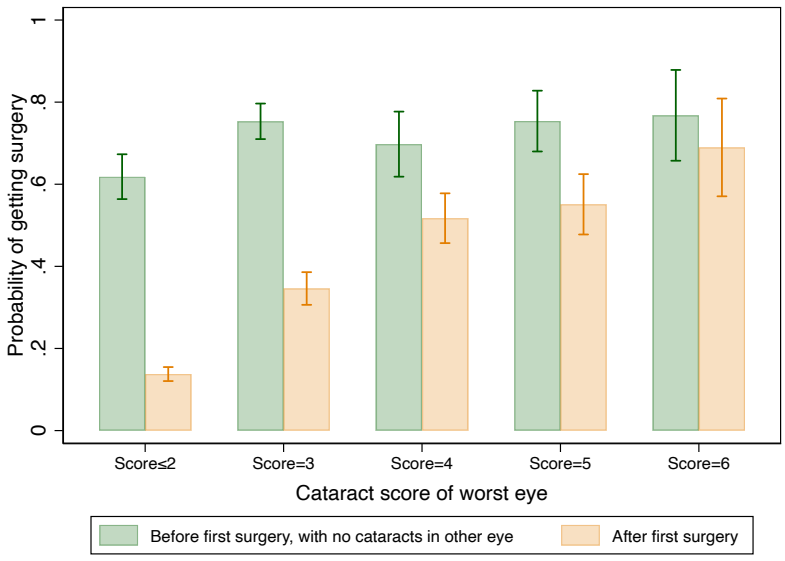


FIGURE 2: Probability of surgery, conditional on one eye having zero cataracts

	Before first surgery		After first surgery	
	Alice	Bob	Alice	Bob
First eye score:	4	6	0	0
Second eye score:	3	3	3	3

To approximate this experiment in the data, we test for differences in take-up rates based on the first-eye cataract score by estimating

$$\mathbb{1}\{\text{Second surgery}_{is}\} = \alpha_s + \beta_s \text{First-eye score}_i + \boldsymbol{\gamma}_s' \boldsymbol{x}_i + \varepsilon_{is},$$

for each patient i and for each second-eye cataract score $s = 1, \ldots, 6$, where $\mathbb{1}\{a\}$ is an indicator of the event a, and $\boldsymbol{x}_i$ is a vector of controls to ensure comparable patients, including gender, age dummies, and insurance status indicators (private, public, none) as a proxy for income.

Table 3 presents the estimation results. We find that worse first-eye scores are associated with lower take-up rates for the second surgery, as indicated by negative and significant coefficients. This suggests that patients who learn more from the first surgery are less likely to return for the second. However, when the second-eye cataract score is 4, this negative association disappears. This may be because cataracts in the second eye are severe enough that the first-eye score no longer matters. Additionally, there is little variation since conditioning on a

TABLE 3: Probability of second surgery as a function of first-eye score

Second eye cataract score is	≤ 2	3	4
First eye cataract score	-0.030***	-0.062**	0.017
	(0.008)	(0.025)	(0.083)
Observations	1,659	576	194
R-squared	0.048	0.104	0.181
Mean dependent variable	0.148	0.439	0.490

Notes: Sample restricted to patients who undergo the first surgery. Robust standard errors in parentheses. Regressions include controls for gender, age dummies, and insurance status fixed effects. p-value test for joint significance across columns is 0.000. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

second-eye score of 4 means the first-eye score can only be 4, 5 or 6.

Note that in this exercise, Alice and Bob may have different preferences for eyesight, which could lead Bob to delay the second surgery until his condition worsens. After all, Alice chose the first surgery with a cataract score of 4, while Bob waited until his score reached 6. Therefore, while we consider the previous result to be suggestive of learning in our data, we need to develop a more structured approach to tease out learning from other selection effects. While each of the previous empirical exercises has limitations, together they suggest the possibility of learning in cataract surgery decisions in our data. Though not definitive, these findings motivate a more structured model of surgery choice to construct a test for learning.

3.1.3 A simple learning model vs full information

We can further illustrate the potential role that the information revealed through undertaking the first surgery plays in patients' decision-making processes with a relatively simple setup. Consider the following oversimplified model of the cataract surgery market. Our actual estimation is based on a much richer and more flexible model, which is formally introduced in section 4. Please note that essentially all the assumptions made in this section will be significantly relaxed in the estimation.

An individual i receives a utility $u_{it} = \alpha_i + \varepsilon_{it}$ from undergoing surgery $t = 1, 2$. The random variables α and ε are independent with mean zero and finite variances. Without loss of generality, we normalize the utility of no surgery to 0. The individual has to decide sequentially to purchase each surgery: that is, the individual decides if she consumes at $t = 1$, receives u_{i1} in that case, and

then decides if she consumes at $t = 2$, and gets u_{i2} in that case. At each t, the realizations ε_{it} are known to the consumer. If we assume *full information*, then the realization α_i is also known to the consumer at $t = 1$, so there is nothing to learn. However, if we assume *learning*, then the realization of α_i is known only after consuming at $t = 1$.[16]

Appendix section C allows for risk aversion (or risk seeking preferences) and also for decreasing marginal utility by penalizing u_{i2}. For simplicity, this section assumes risk neutrality and no such penalization, but results hold if we relax these assumptions.

Assume also that the firm is a monopolist with zero marginal costs that can charge consumers their willingness to pay, p_{it}, at each $t = 1, 2$. That is, the monopolist has the same information set as the consumer. Perfect price discrimination is not essential, but the exposition is easier; appendix section D analyses the non-perfect price discrimination case and reaches similar conclusions.

At $t = 2$, the consumer purchases if and only if $u_{i2} - p_{i2} > 0$, which implies that

$$p_{i2} = \max\{0, \alpha_i + \varepsilon_{i2}\},$$

with both full information or learning, because, by $t = 2$, the consumer knows $\alpha_i + \varepsilon_{i2}$.

At $t = 1$ the consumer purchases if and only if

$$\mathbb{E}_\alpha \left[u_{i1} - p_{i1}\right] + \mathbb{E}_{\alpha,\varepsilon} \left[u_{i2} - p_{i2} | u_{i2} - p_{i2} \geq 0\right] P \left[u_{i2} - p_{i2} \geq 0\right] \geq 0,$$

where the second term represents an option value at $t = 1$, which is present regardless of the informational assumption, but in the equilibrium of this simple model is equal to 0.

Therefore, the price at $t = 1$ depends on the information set at $t = 1$ thus:

$$p_{i1} = \begin{cases} \max\{0, \mathbb{E}_\alpha \left[\alpha_i\right] + \varepsilon_{i1}\} & \text{with learning,} \\ \max\{0, \alpha_i + \varepsilon_{i1}\} & \text{with full information,} \end{cases}$$

because, with full information, the consumer knows α_i at $t = 1$.

All prices are truncated random variables. If we assume normality, then, it can be shown in this simple model that the ratios of variances behave depending

[16] Alternatively, we could assume that "no learning" means that α_{it} for $t = 1, 2$ are never known before the surgery, which implies the consumer does not take them into account in their decision-making process; section B in the appendix explores this alternative and reaches similar conclusions.

on the nature of learning and information:

Lemma 1. Suppose that $\varepsilon \sim \mathcal{N}(0,1)$, $\alpha \sim \mathcal{N}(0,\sigma_\alpha)$, and are independent. Then,

$$\mathbb{V}\left[p_{i2}\right] = \left(\sigma_\alpha^2 + 1\right)\left[1 + 4\phi(0)^2\right],$$
$$\mathbb{V}\left[p_{i1}\right] = \begin{cases} 1 + 4\phi(0)^2 & \text{with learning} \\ \left(\sigma_\alpha^2 + 1\right)\left[1 + 4\phi(0)^2\right] & \text{with full information,} \end{cases}$$

where ϕ is the pdf of a standard normal. In other words, the ratio

$$\frac{\mathbb{V}\left[p_{i2}\right]}{\mathbb{V}\left[p_{i1}\right]} = \begin{cases} 1 + \sigma_\alpha^2 & \text{with learning} \\ 1 & \text{with full information.} \end{cases}$$

Under this simple setup, in the presence of learning, the variance of prices in the second surgery is larger than the variance of prices in the first, which is a testable implication.

We can test this prediction in the data by decomposing the shocks $\varepsilon_{it} \equiv \boldsymbol{\beta}'\boldsymbol{x}_{it} + \widetilde{\varepsilon}_{it}$. Specifically, we can estimate

$$p_{it} = \alpha_i + \boldsymbol{\beta}'\boldsymbol{x}_{it} + \widetilde{\varepsilon}_{it}, \tag{1}$$

where $\boldsymbol{x}_{it}$ is a vector of patient-surgery-specific covariates. For this exercise, we estimate patient fixed-effects regressions and consider time-varying patient and product characteristics as controls: cataract scores for each eye, a time trend, and dummies for type of surgery, for type of lense, for type of insurance, for type of amenities, and for the identity of the sales agent who handled the appointment.

We estimate the variances of prices from the residuals (including the fixed effects). That is,

$$\widehat{\mathbb{V}\left[p_{it}\right]} \equiv \mathbb{V}\left[p_{it} - \widehat{\boldsymbol{\beta}}'\boldsymbol{x}_{it}\right],$$

where the variance on the right-hand side is the sample analogue.

Importantly, before the estimations, we must address selection concerns: not all patients return for a second price quote. Only those that, presumably, had a favorable experience in their first visit return. Therefore, our data lacks price quotes for those patients who chose not to return. We offer further details on how we address selection bias in section 4.6. In a nutshell, we predict those missing prices using a simple machine learning tool with a high degree of accuracy; see section I for further details on the algorithm. As an alternative, we also

estimate (1) using a Heckman two-step control function, in which the excluded instruments are the patient's log age and gender, and results are very similar.[17]

Indeed, using a simple variance-comparison test, we reject the null of equal-sized variances between first and second prices, in favor of $\mathbb{V}[p_{i2}] > \mathbb{V}[p_{i1}]$, with a p-value of 0.0108.[18]

Robustness. While we made strong assumptions behind this simple model, the insights recovered survive relaxing most of them. Consider first the market structure: The firm is not a monopolist, but it does face a downward-sloping demand curve, which implies that prices in the data are indeed informative of willingness-to-pay as required (see also appendix D). Second, we can alternatively estimate (1) without individual fixed effects, which implies that consumers have nothing to learn, as in appendix B, and results are similar. Third, we can allow for partial learning of α_i, which yields a similar model, and similar implications. Fourth, we can estimate (1) separately for each t, which implies a similar model, but also allows for sample sizes at each t to affect the variation of the residuals. We get similar results, even after simulating same sample sizes. Fifth, we can change the order in which we include regressors $\boldsymbol{x}_{it}^k$ and the overall pattern holds. Sixth, we can plot the ratio of price dispersion against R^2, or other measures of fit, and the result holds. And, seventh, we can allow for general risk preferences and diminishing marginal utilities, as in appendix C, and results hold.

Taken together, this descriptive evidence and the theoretical framework suggest that there may be a potential learning channel for these patients. In the model and estimations below, we formalize the different forces that may be at play for patients faced with these decisions.

4 Model

We estimate a (more detailed) model of sequential demand for cataract surgeries, which allows us to capture the amount of information about a second surgery obtained from undergoing a first one. In our model, forward-looking consumers consider whether to undergo surgery in each of their eyes, sequentially, given

[17]The underlying assumption is that patient age and gender do not have a direct impact on price in our setting, given that sales agents use medical recommendations and surgery characteristics to choose from their available menu of prices.

[18]For completeness: we estimate the (residualized) standard deviation of p_{i1} to be 4,417, and that of p_{i2} to be 4,584. If instead we use a Heckman-correction, the p-value is 0.0141.

their current information set and surgery prices. Compared to the simple model of section 3.1.3, here we are able to control for patient and surgery characteristics, which allow us to take into account risk aversion, income effects, diminishing marginal utility, and price endogeneity, among other important confounders.

Before the first surgery, the outcome has an uncertain component.[19] After the first surgery and before the second, uncertainty is (partially) resolved.

As suggested by the survey evidence in section 3.1, patients consider the option value of undergoing a first surgery. Therefore, an important additional feature of the model is that we assume they are forward looking. In sum, patients' demand for the second surgery considers the information revealed through undergoing the first, while their demand for the first surgery also considers the potential information gains regarding a potential second procedure. As such, we assume forward-looking, rational consumers who recognize uncertainty and anticipate that experience will resolve some of it. This approach is benchmarked against an alternative framework in which consumption does not reveal new information. A third, more complex alternative assumes that consumers are initially unaware of the option value (i.e., believing no information will be revealed) but subsequently update their decision-making based on the realized information after the first surgery. Even in this case, a learning process occurs, albeit without full ex-ante recognition of the option value. Hence, we abstract from this third option.[20]

Below, we describe the model in detail, while being explicit about the main underlying assumptions.

As an overview, we model a forward-looking consumer who has to decide whether to undergo surgery in each of her eyes, conditional on her current information set. For the first eye, the outcome of the surgery has an uncertain component, from the consumer's point of view. We assume the consumer has perfect foresight about prices and the rest of the characteristics of the surgery and of herself. The first eye is always chosen by nature as the eye with the worst cataract score. After the consumer has had a surgery for the first eye, the uncertain component is no longer uncertain, because the consumer learns from the experience. The model is compatible with partial learning and flexible risk

[19] The first eye is always chosen by nature as the eye with the worst cataract score.

[20] In the literature, some learning models assume forward-looking consumers, while others rely on myopic consumers who do not account for the value of future information in their current consumption decisions. There is no consensus on which approach better fits the data (Ching, Erdem and Keane, 2017).

preferences. Moreover, the consumer knows that information is revealed after the first surgery, which implies an option value from it.

At each surgery opportunity, consumers face a binary choice of undergoing a surgery or not. That is, except for prices, the characteristics of the surgery are exogenous, and the consumers' consideration sets contain exactly two products: the surgery and the outside option. This assumption is realistic, because the main characteristics of the surgery are determined by the ophthalmologist for medical reasons, and is also consistent with a literature that documents patients adhere to physicians' prescriptions.[21] Moreover, because we study surgery take-up, our main counterfactual only requires us to model realistic substitution patterns towards the outside option.

4.1 Uncertainty about surgery outcomes

Let i index consumers. We denote α_i as the (uncertain) outcome of a surgery for patient i, and assume that α_i is an iid random shock across consumers. The variance of α_i is a measure of the size of uncertainty that consumer i is facing when considering whether to a undergo surgery.

Moreover, we assume that α_i can be decomposed into α_i^k, a knowable component, and α_i^u, an unknowable component. More precisely,

$$\alpha_i \equiv \alpha_i^k + \alpha_i^u.$$

The unknowable component is never revealed to the consumer, and can be thought of as long-term benefits or costs of the surgery which can only be learned over a period longer than the time frame captured in our data.

4.2 Parametrization of Utility

Define the (ex post) utility obtained from undergoing surgery t as

$$u_{it} \equiv \alpha_i + \boldsymbol{\beta}' \boldsymbol{x}_{it} + \varepsilon_{it},$$

for $t = 1, 2$. The outside options are valued at $u_{i0t} = \varepsilon_{i0t}$.[22] Because the consumer does not learn about α_i^u, and does not observe it, we argue that the

[21] See Finkelstein et al. (2021) for the general case of physicians' recommendations, or Johnson and Rehavi (2016) and Gruber and Owings (1996) for the case of cesarean sections.

[22] The reader might think that the outside option for the second surgery is not valued at ε_{i02},

consumer de facto does not experience α_i^u, at least in the relevant time-frame for decision-making. Therefore, we set $\alpha_i^u = 0$, which implies that consumer welfare calculations must be interpreted as the welfare that consumers experience after the surgery, which is relevant for them when they make decisions. Henceforth, we simply set $\alpha_i = \alpha_i^k$.[23]

4.3 Timing

The timing is as follows:

1. Consumer i observes ε_{i01} and ε_{i1}.
2. i decides to operate eye 1 or not.
   (a) If $y_{i1} = 0$, utility is ε_{i01}. End.
   (b) If $y_{i1} = 1$, nature draws α_i from a distribution G_i and ε_{i2}.
3. i observes α_i, ε_{i02}, and ε_{i2}.
4. i decides to operate 2 or not.
   (a) If $y_{i2} = 0$, utility is $\alpha_i + \boldsymbol{\beta}'\boldsymbol{x}_{i1} + \varepsilon_{i1} + \varepsilon_{i02}$. End.
   (b) If $y_{i2} = 1$, utility is $\alpha_i + \boldsymbol{\beta}'\boldsymbol{x}_{i1} + \varepsilon_{i1} + \alpha_i + \boldsymbol{\beta}'\boldsymbol{x}_{i2} + \varepsilon_{i2}$. End.

For reference, the model tree can be found in section E in the appendix.

4.4 Demands

Demands can be derived by backward induction. The consumer decides to get the second surgery if and only if

$$u_{i2} = \alpha_i + \boldsymbol{\beta}'\boldsymbol{x}_{i2} + \varepsilon_{i2} > u_{i02} = \varepsilon_{i02}.$$

Therefore, conditional on the first surgery, the demand for the second surgery

because the patient might "carry" α_i^k with her to another provider. However, we assume α_i^k subsumes everything related to ocular health information, and $\alpha_i^u + \varepsilon_{i2}$ subsumes everything else. Therefore, we can interpret $\alpha_i^u + \varepsilon_{i2}$ as including remaining uncertainty from any of the outside options. This assumption is standard in the literature.

[23] Claude Monet's α. Monet is not a representative patient for our sample, but is the most famous cataracts patient, and his case is illustrative. Figure 12 in appendix F shows how Monet perceived the world before and after cataracts. Famously, Monet was advised by friends, family, and physicians to get a cataract surgery, but he was hesitant. After, presumably, their advice pushed his expected marginal utility into a positive sign, Monet underwent the surgery. Then, Monet went back and destroyed some paintings he had created while his vision was impaired. That is, Monet had a positive realization of the α shock, which made him regret some of his work ex post. See Gruener (2015).

is

$$P\left[y_{i2}=1|y_{i1}=1\right]=P\left[u_{i2}-u_{i02}>0\right],$$

where we drop the conditional statement, because of independence.

Then, the expected marginal utility of the second surgery is

$$\mathbb{E}\left[u_{i2}-u_{i02}\right]=\mathbb{E}\left[u_{i2}-u_{i02}|u_{i2}-u_{i02}>0\right]P\left[u_{i2}-u_{i02}>0\right]$$

where the expectations are with respect to $\alpha_i+\varepsilon_{i2}-\varepsilon_{i02}$.

Therefore, before the first surgery, after ε_{i1} and ε_{i01} are known, but before α_i, ε_{i2}, and ε_{i02} are known, the expected marginal utility from the first surgery is

$$\mathbb{E}_{\alpha_i}[u_{i1}-u_{i01}+\mathbb{E}\left[u_{i2}-u_{i02}\right]]=\mathbb{E}_{\alpha_i}[\alpha_i]+\boldsymbol{\beta}'\boldsymbol{x}_{i1}+\varepsilon_{i1}-\varepsilon_{i01}+\mathbb{E}\left[u_{i2}-u_{i02}\right],$$

where the expectation with respect to α_i has been subscripted, and $\mathbb{E}\left[u_{i2}-u_{i02}\right]$ represents an option value conditional on undergoing the first surgery. Note that the option value is always present for the demand of the first surgery regardless of the informational assumptions (that is, even with perfect information of α_i): because of the sequential nature of consumption, if p_2 goes up, the option value goes down, so p_1 must also go down, which makes them complements. In other words, it is not the information structure which makes surgeries (even more) complements, but the sequence of events.

The consumer chooses the first surgery if and only if

$$\mathbb{E}_{\alpha_i}[u_{i1}-u_{i01}+\mathbb{E}\left[u_{i2}-u_{i02}\right]]>0.$$

Therefore, the demand for the first surgery is

$$P\left[y_{i1}=1\right]=P\left[\mathbb{E}_{\alpha_i}[u_{i1}-u_{i01}+\mathbb{E}\left[u_{i2}-u_{i02}\right]]>0\right].$$

4.5 Log-likelihood

We begin by assuming distributions for the shocks. A common assumption is Type-1 extreme-valued ε-shocks, which yield a mixed logit model if we assume a normal distribution for α. However, because we have a binary choice model, we see no advantage of a mixed logit vis-à-vis normally distributed ε-shocks, but, if we assume normality, we can solve analytically for the equilibrium of the model.

We make the following simplifying assumption.

Assumption 1. $\varepsilon_{i1} - \varepsilon_{i01}$ and $\varepsilon_{i2} - \varepsilon_{i02}$ are iid $\mathcal{N}(0,1)$, and α_i are iid $\mathcal{N}(\mu_{\alpha,i}, \sigma_{\alpha,i})$.

Under assumption 1, we obtain

$$\mathbb{E}\left[u_{i2} - u_{i02}\right] = \left[\boldsymbol{\beta}'\boldsymbol{x}_{i2} + \mu_{\alpha,i} + \sqrt{1+\sigma_{\alpha,i}^2}\lambda\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2} + \mu_{\alpha,i}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right)\right]\Phi\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2} + \mu_{\alpha,i}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right)$$

$$P\left[y_{i1} = 1\right] = \Phi\left[\mu_{\alpha,i} + \boldsymbol{\beta}'\boldsymbol{x}_{i1} + \mathbb{E}\left[u_{i2} - u_{i02}\right]\right] \tag{2}$$

and

$$P\left[y_{i2} = 1 | y_{i1} = 1\right] = \Phi\left(\frac{\mu_{\alpha,i} + \boldsymbol{\beta}'\boldsymbol{x}_{i2}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right), \tag{3}$$

where λ is the inverse Mills ratio: $\lambda(z) \equiv \phi(z)/\Phi(z)$, with ϕ and Φ the pdf and cdf of a standard normal. Section E in the appendix shows the derivations for these expressions.

From (2) and (3), we can see that if we assume $\mu_{\alpha,i} = \mu_\alpha \ \forall i$, then, μ_α is not separately identified from the constant. Therefore, we assume:

Assumption 2. $\forall i,\ \mu_{\alpha,i} = 0$.

Note assumption 2 is consistent with ex ante informed consumers who can correctly anticipate their mean utility level from surgeries.

Finally, let $s_{i1} \equiv P\left[y_{i1} = 1\right]$ and $s_{i2} \equiv P\left[y_{i2} = 1 | y_{i1} = 1\right]$, and recall our parametrization $\sigma_{\alpha,i} \equiv \exp(\boldsymbol{\theta}'\boldsymbol{w}_i)$. Therefore, the log-likelihood of the data becomes

$$\ell = \sum_{i=1}^{N} y_{i1}\log s_{i1} + (1-y_{i1})\log(1-s_{i1}) + y_{i1}y_{i2}\log s_{i2} + y_{i1}(1-y_{i2})\log(1-s_{i2}),$$

which is maximized for $(\boldsymbol{\beta}, \boldsymbol{\theta})$.

4.6 Identification and estimation

To estimate the model, we need to identify two key parameters: the price coefficients in each surgery and the uncertainty shock. We identify price coefficients using sales targets as instrumental variables. To identify the uncertainty shock, we leverage differences in a patient's behavior between their first and second

surgery. Under the assumption that the price coefficient is accurately identified, any decrease in a patient's price elasticity for the second surgery must be attributed to the impact of the uncertainty shock, $\sigma_{\alpha,i}$, on the price coefficient for the second surgery.

We address important threats to identification; namely, price endogeneity, selection bias, and confounders such as risk aversion, diminishing marginal returns, and income effects. We also discuss the identification argument for the variance of uncertainty shocks.

To be specific, the covariate vector $\boldsymbol{x}_{it}$ includes: log prices, log age, gender, type of insurance as a proxy for income, a risk aversion proxy, cataract scores of each eye (which we term as minimum and maximum scores, instead of using left and right eye), dummies for type of surgery, type of intra-ocular lens, type of amenities, and sales-agent fixed effects. The covariate vector $\boldsymbol{w}_i$ includes: log age, gender, cataract scores, and ocular health measures.

Price endogeneity. To identify the price coefficient, we use sales targets variables as instruments in a control function strategy described in detail in the appendix section G (Petrin and Train, 2010). Specifically, we use the agent-specific 15-day rolling-average log price up to the moment the sales agent was talking with the consumer. Intuitively, on a slow day or week, the agent might decrease offered prices to close a sale, and on a good day or week, the agent might offer high prices. However, how slow a day or week has been for the agents should not directly affect the current consumer. Specifically, the patient-specific demand shocks are uncorrelated with how far or close an agent is to her target on a given day or week. For instance, the patient's support system for post-operative care should be orthogonal to whether a particular day faced an unusually low demand. Generally, as long as the demand shocks are independent across patients, the exclusion restriction will hold.[24] Moreover, we include agent fixed effects in our estimation to control for agent-specific price biases or persuasion techniques. To alleviate the incidental parameter problem due to agent fixed effects, we keep only those agents with 8 or more quotes in the data (Greene, 2004), which is more than 98% of quotes.[25] Appendix H shows the results of the first stage.

[24] One possible violation of this assumption would be if, on a given day, there was unusually high traffic or roadblocks that made access to the clinic more difficult. However, it does not seem that many patients arrive by car, possibly due to the central location of the main clinic (a mix of both availability of public transportation and high parking costs).

[25] Results are robust if we use the following alternative instruments: sales agent's previously offered price, if the previous price quote was taken, the percentage of operations sold during the week.

Risk aversion. As a potential confounder, it may be that patients who return are a selected sample of unobservably more risk averse individuals (the same reasoning applies to sicker or higher income patients). There are essentially two types of risk in this setting: First, there is a surgery-specific risk due to potential complications, such as infections; second, there is the uncertainty shock α_i, which represents an individual-specific surprise outcome. While the second type of risk is explicitly modeled, we need to address the first one, because, otherwise, it would be absorbed by the ε_{it} shocks. Our data allow us to measure a proxy for risk aversion, under the reasonable assumption that patients who engage in more visits before obtaining a price quote, but *after* their cataract diagnosis, are more risk-averse than patients who do so in fewer visits, all else equal (including observable health characteristics). Hence, we control for (log) visits per price quote, which offers variation at the surgery- and individual-level. Moreover, our proxy for risk aversion subsumes other kinds of unobserved heterogeneity and risks, which we exclude from α_i. Note, however, that the uncertainty in α_i might still subsume some possible forms of risk aversion; i.e., if consumers are indeed risk averse, the α_i should really be (akin to) $\alpha_i + 1 - e^{\rho\alpha_i}$ (similar to appendix C), and the estimated distribution of α_i is that of $\alpha_i + 1 - e^{\rho\alpha_i}$. If this were the case, then there is a relatively minor misspecification in our approach: around risk neutrality ($\rho = 0$), a first-order approximation yields $(1 - \rho)\alpha_i$ instead of just α_i.

Decreasing marginal returns. We allow for consumers to have different marginal valuations of surgeries as a function of health characteristics. For instance, the first surgery might be more valuable, because the patient might recover a larger improvement in their vision from it. The second surgery might be less valuable, because the patient already has an improved vision. We provide flexibility for valuations to go either way by controlling for cataract scores of both eyes in both surgery demands, as well as for other patient and surgery characteristics.

Income effects. Because surgeries are relatively expensive, even at our partner firm, we might have income effects, where high-income consumers are less elastic, for example. In Mexico, high income is highly correlated with health insurance coverage (ENSANUT, 2020). Therefore, as a proxy for income we use type of insurance, which is observable in our data. Alternatively, we also observe zip codes, but we favor insurance type, because results are similar and less prone to incidental parameters problems.

Size of uncertainty shocks. We further parameterize

$$\sigma_{\alpha,i} \equiv \exp(\boldsymbol{\theta}' \boldsymbol{w}_i),$$

where w_i are some time-invariant patient characteristics, and $\boldsymbol{\theta}$ is a vector of parameters to be estimated. Intuitively, the magnitude of shocks, $\sigma_{\alpha,i}$, is identified from the discrepancies between the covariates' effects from the first operation and the second one. The coefficients $\boldsymbol{\theta}$ are identified from the correlations between the magnitudes of $\sigma_{\alpha,i}$ and covariates $\boldsymbol{w}_i$. More formally, we identify $\sigma_{\alpha,i}$ given that ε-shocks have the same variance: Conditional on observables, if the decision to undergo either surgery is the same between surgeries, except for the information set, then we have identification. In this context, we find this assumption to be reasonable, because we observe the major components of the decision; namely, prices, surgery characteristics, demographic characteristics (e.g., type of insurance coverage), a proxy for risk aversion, and ocular health measures. Moreover, the utility specification is flexible enough to account for a lower marginal utility for the second operation and for the role of cataract severity in each of both eyes. Finally, this assumption is very common in the dynamic discrete choice models literature (Arcidiacono and Ellickson, 2011).

Selection bias. The outstanding issue is that of selection into our sample of second surgeries. Consumers who return for a consultation about the second surgery presumably had a positive shock from the first one. Consumers who never return are not in the data and the counterfactual prices of a second surgery are unobservable to us. If we ignore this fact, we might overestimate the benefit from the surgery.

Therefore, when missing, we predict the (log) price that a consumer would have had if she came back for a second consultation, as a function of characteristics of patients, surgeries, sales agents, optometrists, and ophthalmologists. We use a simple machine learning technique, LASSO, to predict prices; we find LASSO outperforms a linear regression in this setting, as measured by the mean prediction error, and achieves an R^2 of .72. We match the distribution of predicted prices to the empirical distribution of non-missing prices, as shown in Figure 13, which can be found in the appendix section I, along with further details of the algorithm.

Estimation. Estimation is performed in two steps. We first construct a control function using the price instruments described above and explained with detail in appendix G. Then, we add the control function as an extra regressor, and we

perform a maximum likelihood estimation parameterized as above. We run 500 (panel) bootstraps of the whole process to calculate standard errors.

5 Results

The estimation is carried out on our sample of patients whose initial contact with our partner firm occurred in 2018 and had at least one cataract-related visit. We follow the procedure described in Section 4.6 for estimating parameters via maximum likelihood over 500 bootstrap repetitions. Section 4.6 also discusses important threats to identification.

Table 4 shows our estimated elasticities and associated standard errors clustered at the patient level. The top panel shows estimates for the indicator variable for whether the patient takes up the surgery, and the bottom panel corresponds to our uncertainty shock parameter $\sigma_{\alpha,i}$. Different columns offer different specifications of the σ equation and some add a control function to address the endogeneity of prices based on the 15-day rolling average of log prices offered by the agent as an instrument. Importantly, one of such controls corresponds to sales agent fixed effects, which addresses any unobservable heterogeneity in sales tactics, such as persuasion or announcing ad hoc prices.

All columns include the following control variables: log of the patient's age, gender, cataract scores for each eye (labeled as the minimum and maximum scores, regardless of whether they correspond to the left or right eye), a proxy for risk aversion (measured as the log of the number of visits per price quote after the diagnosis), characteristics of the surgery, a proxy for income (measured as type of patient's health insurance), and sales agents fixed effects. For clarity, only coefficients of interest are reported.

As a benchmark, column 1 shows a standard IV-probit, where we do not allow for uncertainty shocks, α, nor we allow for an option value of the second surgery. Column 2 allows for uncertainty shocks, but does not address endogenous prices, and yields very elastic demand curves. Column 3 addresses price endogeneity.

Column 4 shows our preferred specification, which includes both the control variables in the σ equation and a control function approach. Intuitively, patient heterogeneity matters for both the decision to get the surgery and the uncertainty parameter. As expected, we estimate a negative and significant effect of prices

[26]The unreported estimates for surgery characteristics are significant. Estimates for log age are not as small, but are all insignificant. These point estimates seem to suggest that perhaps older people are more likely to take up surgery.

TABLE 4: DEMAND ESTIMATIONS

DEP VAR: Operates$_{it}$	(1)	(2)	(3)	(4)
log price	−1.43	−0.97	−2.12	−2.13
	(0.439)	(0.078)	(0.671)	(0.701)
log Age	0.00	−0.09	−0.05	−0.09
	(0.058)	(0.085)	(0.094)	(0.201)
Female	−0.04	−0.09	−0.10	−0.05
	(0.028)	(0.041)	(0.040)	(0.108)
Min cataract score	0.02	0.07	0.08	0.14
	(0.010)	(0.015)	(0.019)	(0.058)
Max cataract score	0.07	0.03	0.02	0.00
	(0.012)	(0.017)	(0.017)	(0.057)
Risk aversion proxy	0.68	1.13	1.01	0.94
	(0.048)	(0.032)	(0.128)	(0.132)
DEP VAR: $\sigma_{\alpha,i}$				
log Age				−0.05
				(0.260)
Female				−0.15
				(0.179)
Min cataract score				−0.14
				(0.104)
Max cataract score				−0.02
				(0.104)
RE health score				0.00
				(0.087)
LE health score				−0.01
				(0.094)
cons		3.91	3.83	3.46
		(0.103)	(0.134)	(1.202)
ELASTICITIES				
ALL OPS	-1.51	-1.85	-3.80	-3.29
FIRST OPS	-1.52	-2.57	-5.26	-4.51
SECOND OPS	-1.49	-1.11	-2.31	-2.04
OTHER CONTROLS	YES	YES	YES	YES
CONTROLS ($\sigma_{\alpha,i}$)	NO	NO	NO	YES
CONTROL FUNCTION	YES	NO	YES	YES
MPE	0.39	0.39	0.37	0.36
R^2_p	0.02	0.05	0.08	0.11
FIRST-STAGE IV'S F	29.91		29.91	29.91
PATIENTS	3,822	3,822	3,822	3,822
QUOTES	7,627	7,627	7,627	7,627

Notes: Bootstrapped standard errors clustered at individual level with 100 repetitions. Min and max cataract scores are a patient's eyes minimum and maximum score in a 0 to 6 scale. Risk aversion proxy is (log) number of visits to obtain a price quote. Control functions for prices are constructed with the 15-day rolling-average of log prices up to the moment as an instrument (Petrin and Train, 2010). Other controls include: a constant, sales agents fixed effects, surgery characteristics, and type of insurance as proxy for income. In the $\sigma_{\alpha,i}$ equation, eye health scores count the number of comorbidities (ampliopia, anisometropia, astigmatism, myopia, presbyopia, hypermetropia, and emmetropia) present in each eye. MPE stands for mean prediction error. R^2_p stands for pseudo-R^2, constructed as described in footnote 27. Price elasticities with respect to unconditional demands: $P[y_{it} = 1], t = 1, 2$.

on surgery take up. For our patient characteristics, we find insignificant effects that are also very close to zero.[26]

As for the information shock, we find that older people, women, and patients with worse scores experience lower uncertainty from the first surgery. The medical literature has documented that take-up of cataract surgery in low- and middle-income countries is consistently lower for women than men, but has been unable to provide a convincing explanation for this differential (Mercer, Lyons and Bassett, 2019; Briesen et al., 2010). In our estimation, female gender has a small and insignificant effect on demand (and if anything, is negative), but is related to lower uncertainty (although also statistically insignificant), which in turn increases the elasticity for the second surgery.

All columns support the existence of an information shock, where the null hypothesis is no information shock. We also find that a model that allows for heterogeneity in $\sigma_{\alpha,i}$ fits the data slightly better than one without heterogeneity, as measured by the mean prediction error and a pseudo-R^2.[27] Indeed, in the appendix J, we perform our estimations in a training set, consisting of a random sample of 50% of the data, and we test our estimations in the hold-out sample of the remainig 50%. We find our preferred specification outperforms the rest, and is not simply overfitting the data.[28]

Our overall elasticities are in line with what one might expect: because patients mainly pay out-of-pocket and this is an elective procedure, we find elastic demands. These estimates are broadly consistent with other Latin American settings. For instance, Duarte (2012) exploits variation in Chile in public sector price caps that affects private insurance plans, finding that elasticities for elective procedures range from -0.3 to -2.1. In the US, researchers analyzing the seminal RAND health insurance experiment obtained very inelastic demands (Newhouse, 1993; Lurie et al., 1989). However, some of these elasticities have been revisited in various ways (see Aron-Dine, Einav and Finkelstein (2013) for a broad discussion). For instance, Kowalski (2016) uses a censored quantile instrumental variable estimator that leads to elasticities for medical care between -0.8 and -1.5, an order of magnitude larger than the original RAND experiment

[27] In this paper, R^2_p is constructed as (number of correct predictions - number of most frequent outcome) / (number of outcomes - number of most frequent outcome).

[28] We have also performed a Vuong (1989)'s nonnested model test, where we constructed a likelihood ratio test of our model relative to a model in which we have no learning nor dynamics, equivalent to column 1 in Table 4. We find that we can reject the null in favor of nonnested models with a t-stat of 46.6. Although these statistical tests do not specify what consumers are learning about, they do indicate that patients learn something with the first surgery.

elasticities.

We find that elasticities are consistently larger for the first operation compared to the second. In our preferred specification, we find that a 10% increase in the price leads to an overall decline of 32.9% in the probability of getting cataract surgery. However, for a 10% increase in the price of the first surgery, take-up goes down by 45.1%, while the same percentage change for the second surgery only leads to a decline of 20.4% in the probability of take-up.

We also note that our risk aversion proxy is insignificant, except when prices are not instrumented. This makes sense because, if price effects are not properly identified, then differential elasticities between operations have to be (erroneously) explained through differential risk attitudes. Indeed, to rationalize observed demands, column 2 finds that people who are more risk averse are also more willing to undergo surgery, which yields unrealistic price elasticities.

Figure 3 plots the distribution of our estimated uncertainty shock parameter. We find considerable heterogeneity, with bimodality (due to gender dummies), and a relatively long right tail (for clarity, the plot winsorizes the distribution at the 95th percentile). To put it in perspective, the standard deviation of the demand shocks, ε, is equal to 1. We predict that, on average, the surprise component of the surgery is 5.8 times as large as the unobservable shocks, ε. This suggests that the option value from revealed information after the first surgery might be quite large.[29]

Lastly, we measure consumer surplus with the ex post utility from getting surgeries, and we transform consumer surplus into dollar terms by considering the marginal utility of income implied by the model. Specifically, given α_i, and ε-shocks, we compute the ex post surplus as

$$CS_i(\alpha_i, \varepsilon_{i1}, \varepsilon_{i2}, \varepsilon_{i01}, \varepsilon_{i02}) \equiv \mathbb{1}\left\{u_{i1}^m + \mathbb{E}\left[u_{i2}^m\right] > 0\right\} \left(\frac{u_{i1}^m}{\frac{\partial u_{i1}^m}{\partial p_{i1}}} + \mathbb{1}\left\{u_{i2}^m > 0\right\} \frac{u_{i2}^m}{\frac{\partial u_{i2}^m}{\partial p_{i2}}} \right),$$

where $u_{it}^m \equiv u_{it} - u_{i0t}$ are the marginal utilities, and where we simulate 500 vectors of shocks, and average across them to find our estimate of consumer surplus, CS_i.[30]

Figure 4 presents the distribution of the estimated ex post consumer surplus

[29] To be sure, our estimations include a rich set of controls, which decrease the remaining variation due to the unobservable demand shocks.

[30] We obtain qualitatively similar results with ex ante computations of consumer surplus.

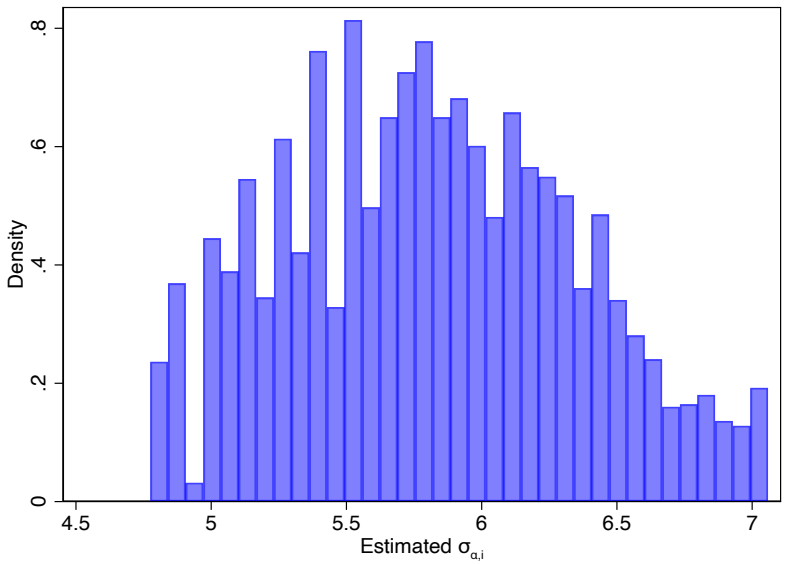


FIGURE 3: Distribution of $\widehat{\sigma}_{\alpha,i}$

Note: Distribution winsorised at 1st and 99th percentiles.

in US dollars, which is on average 288 USD.[31] The distribution presents a high dispersion. We therefore winsorize this plot at the 1st and 99th percentiles. A small percentage of patients have a negative estimated surplus, because they received a negative shock after the first surgery and did not exercise their option for the second surgery. As a reminder, the average price of cataract surgery at this provider is around 13,000 pesos or 700 USD (Table 2).

The interpretation of the consumer surplus is anchored by the outside option. That is, we estimate the marginal surplus with respect to the outside option. Then, the conservative interpretation is that we find a lower bound on the real consumer surplus from the surgery. However, as confirmed through a phone survey, for the vast majority of patients, the outside option is simply not getting a surgery. Therefore, our estimated surplus is close to the real surgery surplus for these patients.

6 Counterfactuals

From the firm's perspective, a key question is whether managers can implement strategies that account for uncertainty in outcomes and the option value of the

[31] During this period, 1 USD = 19.22 MXN.

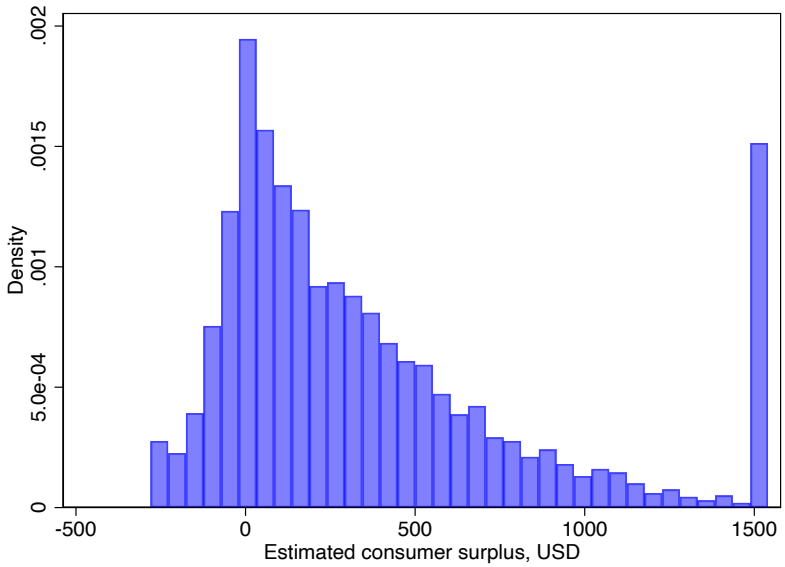


FIGURE 4: Distribution of estimated ex post consumer surplus

Note: Distribution winsorised at 1st and 99th percentiles.

first interaction to improve performance. While increasing revenue is a natural objective, particularly if surgeries are under-consumed, this firm also targets low-income patients and has a mission to expand healthcare access. Enhancing its reputation by increasing surgeries among its target population may therefore be another priority. In this section, we simulate counterfactual scenarios to explore how the firm can balance revenue generation with expanding access to care for low-income patients.

With our estimated parameters, we simulate two counterfactual policies. First, we quantify the welfare costs or gains associated with unexpected outcomes in a counterfactual scenario where uncertainty is resolved, such as through marketing campaigns by the firm, a central planner, or another stakeholder. The medical literature has explored similar interventions (Mailu et al., 2020).

Second, given the firm's focus on serving low-income patients, we examine whether revenue-neutral price adjustments can increase the number of surgeries performed. This provides a benchmark for price policies that balance revenue

[32]The reader might ask about an obvious counterfactual: tie-ins, that is, bundling both surgeries from the start at a single price. However, because the consumer is uncertain about outcomes, the consumer's dominant strategy is to consume sequentially, which would go against tie-ins. On the other hand, the firm would have to drop the price significantly to convince consumers to purchase the bundle, which is not optimal for the firm. Moreover, the authors have found no firm in this market which offers this option.

generation with expanding healthcare access. Specifically, we consider a budget-neutral subsidy for the first surgery, offset by an increase in the second. This counterfactual is particularly informative, as it increases take-up through an efficient mechanism.

Throughout we analyze the heterogeneous effects of these policies.[32]

6.1 Quantifying uncertainty and persuasion

In this section, we consider a counterfactual unveiling of α_i. This scenario could be interpreted as a hypothetical persuasion or educational campaign, where "champion" patients inform potential consumers about their successful results. These types of interventions have been implemented in a variety of settings and have been experimentally evaluated by the medical literature (Mailu et al., 2020). For instance, a champion c might reveal their outcome $\alpha_c = \sigma_i/2$ to a potential consumer i. The potential consumer i might believe totally or partially in this information. For simplicity, we assume patients completely believe the champion's announcement, but the model can readily incorporate partial persuasion. Moreover, we abstract from the costs of this marketing campaign, which may or may not be borne by the firm.

That is, given a *revealed* α, we compute the ex post consumer surplus with known α as:

$$CS_i^{\alpha}(\varepsilon_{i1}, \varepsilon_{i2}, \varepsilon_{i01}, \varepsilon_{i02}) \equiv \mathbb{1}\left\{u_{i1}^m + \mathbb{E}\left[u_{i2}^m\right] > 0\right\} \left(\frac{u_{i1}^m}{\frac{\partial u_{i1}^m}{\partial p_{i1}}} + \mathbb{1}\left\{u_{i2}^m > 0\right\} \frac{u_{i2}^m}{\frac{\partial u_{i2}^m}{\partial p_{i2}}} \right)$$

where we set $\sigma_{\alpha,i} = 0$, and expectations are only with respect to demand shocks. In other words, if we set $\alpha_i = \alpha$, patients believe the shock to be α with probability 1, and the only remaining uncertainty at the time of the first surgery are the second surgery demand shocks. Then, a reduced option value remains. Again, we simulate a vector of shocks and take an average to find CS_i^{α}.

We offer a range of possible outcomes based on the level of α_i that potential consumers might believe, including $\alpha_i = 0$, which would quantify the value of uncertainty. In this case, when consumers completely believe their ex post utility will match their expected utility, we find a negative estimation for consumer surplus of around -230 USD.

Figure 5 shows how consumer surplus changes from the status quo to a revealed information counterfactual. As expected, as the revealed shock increases,

consumer surplus increases (which also implies increases in the firm's revenue). Perhaps surprisingly, at $\alpha_i = 0$, this change is negative, which implies uncertainty is valuable for consumers. Intuitively, patients value uncertainty, because a bad draw from the surprise distribution can be mitigated by operating just once, but a good draw can be amplified by operating twice.[33]

Therefore, for consumers to value certainty, the revealed information needs to be credible and sizable. If the cost of persuasive advertising is increasing in the "quality" or type of information provided, this approach may not be ideal for a firm that bears the full expense of the campaign. However, if an NGO or the government funds the campaign, the firm may benefit, though this may not be the most efficient allocation of resources for those other stakeholders.

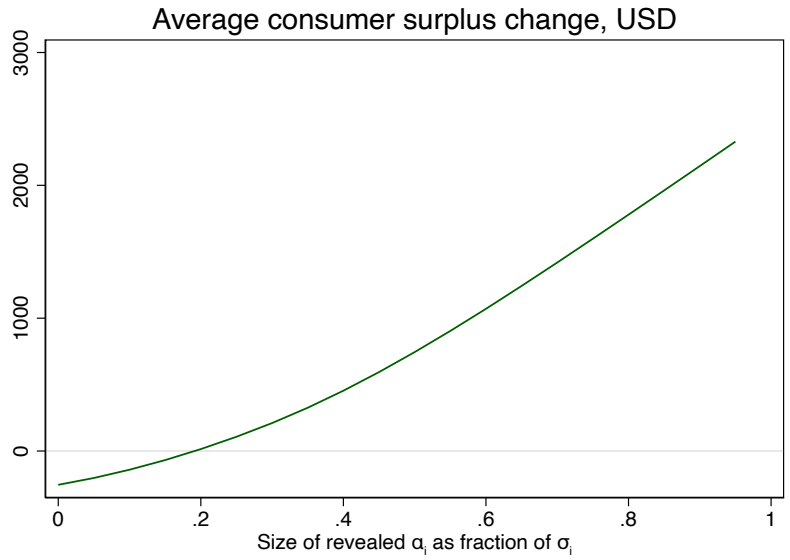


FIGURE 5: Quantifying uncertainty and champions policy

6.2 Revenue-neutral price changes

Now we consider how the firm can increase surgeries through price changes, while maintaining a neutral budget. Beyond managers, policymakers—government, NGOs, or other third-parties—may be interested in such interventions, because

[33] In equilibrium, the firm would react to a counterfactual resolution of uncertainty by changing prices. For instance, if the option value is high, then prices for the first surgery would be high as well, because willingness-to-pay for the first surgery is high. Without uncertainty, the option value drops, because the consumer has less to learn.

take-up might increase through simple price changes. Throughout these counterfactuals we consider ex post estimations of consumer surplus, demand, and revenue, and we focus on revenue-neutral policies as a benchmark that balances the firm's revenue maximizing objective with its goal of increasing healthcare access.[34]

For a forward-looking consumer, price hikes on the second surgery reduce the option value. However, price reductions on the first surgery increase demand through both the first and second surgery. Indeed, if the expected demand of consumer i is $D_i \equiv s_{i1} + s_{i1}s_{i2}$, then,

$$\frac{\partial D_i}{\partial p_{i1}} \frac{p_{i1}}{D_i} = \frac{\partial s_{i1}}{\partial p_{i1}} \frac{p_{i1}}{s_{i1}},$$

where we see the elasticity for total demand is the elasticity for the first surgery. Therefore, it is not obvious if a demand-increasing, budget-neutral price change can be found. But, if the demand for the second surgery is less elastic, we might find a pricing schedule where welfare increases through cross-subsidizing.

Figure 6 shows a counterfactual discount for the first surgery accompanied by an offsetting price increase in the second surgery. In this exercise, if p_1 decreases by x%, then p_2 increases by x%. The firm is roughly indifferent from a revenue perspective, but take-up increases, which may be another objective for the firm, and consumers are significantly better off.

Figure 7 shows an asymmetric price change: if p_1 decreases by x%, then p_2 increases by $.5x$%. In this case we find both the firm's revenues and the consumer surplus are even higher. In particular, take-up increases to a greater extent.

In both cases, new consumers undergo the surgery. Figure 8 shows a breakdown of changes in demand by first and second surgeries. We can see the increase in the extensive margin is mainly due to an increase in the take-up of the first surgery.

In light of these results, a natural question may be why the firm has not implemented these pricing schedules yet. We posit two hypotheses, but are unable to test for them within our model. First, the firm is presumably already near an optimum: In fact, our estimations imply the firm profits about 10% of the surgery's price on average, which coincides with their declared business model. With slim profit margins, the firm might not be able to justify the potential backlash from an ill-executed price hike in the second surgery, including a negative impact

[34]We assume constant marginal costs, which is sensible in this context. The firm has relatively low marginal costs from surgeries, mainly because they pay surgeons by day, not by surgery.

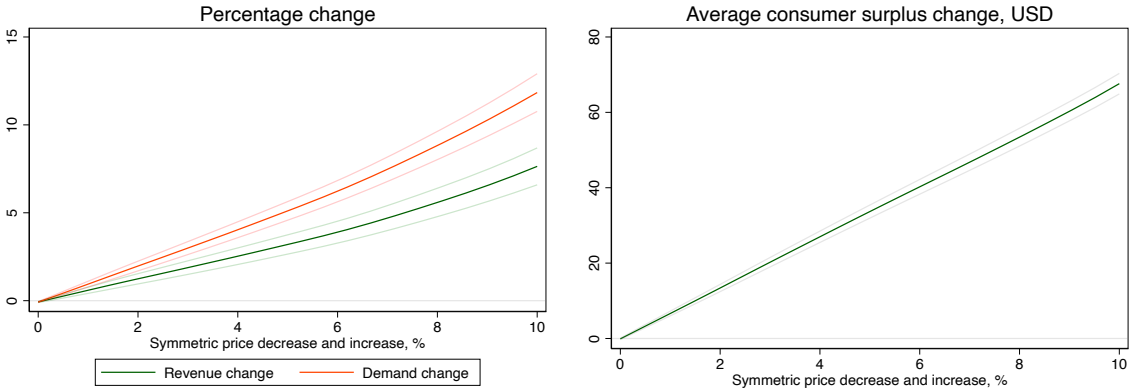


FIGURE 6: Symmetric, revenue-neutral change in pricing policy

Note: Counterfactual decrease in first surgery prices in same percentage as increase of second surgery prices. For example, if p_1 drops in 1%, then p_2 increases in 1%. Lighter lines correspond to 95% confidence intervals.

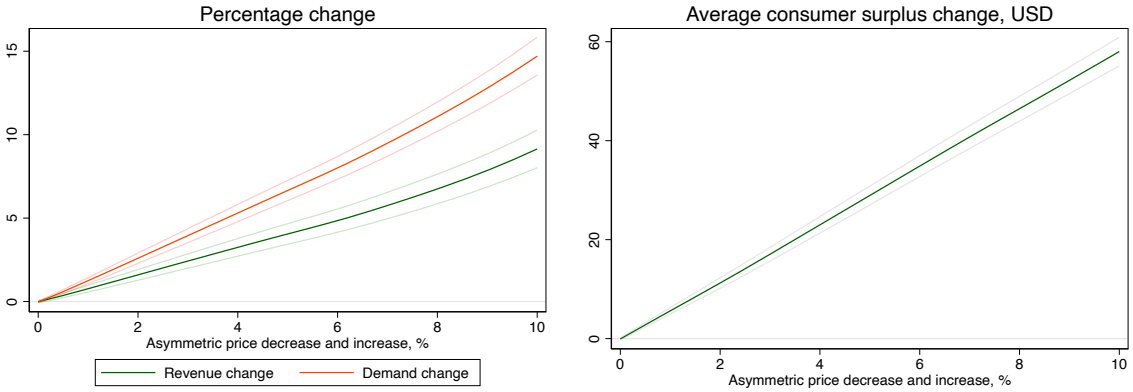


FIGURE 7: Asymmetric, revenue-neutral change in pricing policy

Note: Counterfactual increase of second surgery prices is .5 times the decrease of first surgery prices. For example, if p_1 drops in 1%, then p_2 increases in .5%. Lighter lines correspond to 95% confidence intervals.

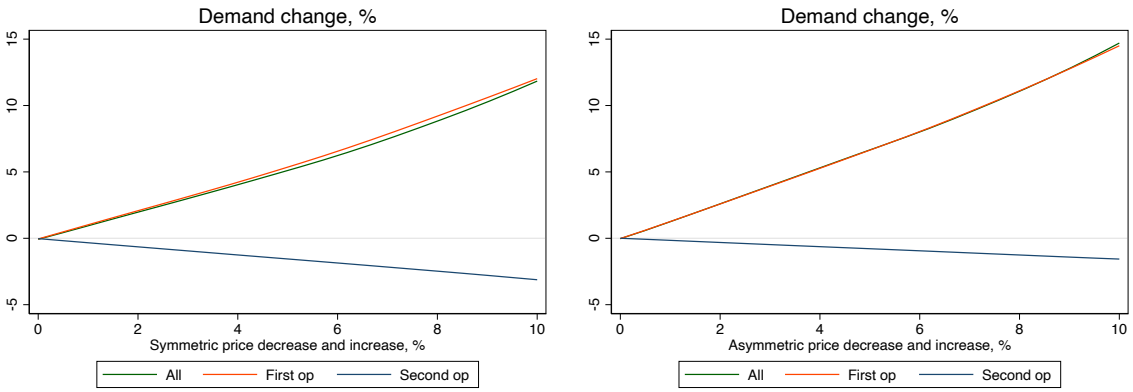


FIGURE 8: Demand changes by surgery, revenue-neutral change in pricing policy

Note: For the asymmetric price change, counterfactual increase of second surgery prices is .5 times the decrease of first surgery prices. For example, if p_1 drops in 1%, then p_2 increases in .5%.

on reputation. Although our proposed price changes can positively impact the firm's reputation through greater healthcare provision, some consumers may also find fault with a higher price on the second surgery. Second, the firm might find it hard to implement this type of price discrimination in practice. Indeed, the industry standard in our setting appears to be independent, sequential pricing.

Therefore, we believe there is room for policy. For instance, the government or an NGO might intervene by offering vouchers or cross-subsidies. These third-parties should find the firm is willing to cooperate, because the firm remains roughly indifferent in terms of revenue, while take-up increases.

7 Concluding remarks

In many experience goods markets, the number of potential repeated purchases might be small and limited. Such is often the case with durable goods or elective healthcare treatment procedures. Given this feature, classic insights about the value of learning on demand and the efficacy of strategies that may increase initial take-up may not necessarily hold true. With these limits on repeated interactions, firms may not be able to successfully adapt their pricing and advertising strategies in order to increase market share, and potential customers may be constrained in their ability to exploit these consumption dynamics. Understanding these factors as well as the size and value of the initial uncertainty is therefore important for quantifying welfare.

To shed light on these issues, we focus on modeling and estimating demand

for cataract surgeries. Exploiting a rich dataset from a large private provider in Mexico City and leveraging sales targets set by the firm for its sales agents, we identify structural demand parameters detailing price elasticities for each of two potential surgeries as well as the value of the uncertainty shock. Our results show that the estimated elasticities are larger for the first surgery and that there is considerable heterogeneity in the idiosyncratic uncertainty parameter. This suggests that the option value of the first surgery is large. We also find heterogeneity in our measure of consumer surplus.

With our parameters, we examine whether managers can implement strategies that incorporate outcome uncertainty and the first interaction's option value while balancing revenue generation and healthcare access. The first set of simulations considers informational interventions akin to persuasive advertising, where the objective is to resolve the uncertainty. We find that reducing uncertainty is only welfare-improving if the firm is able to convince patients of a very positive outcome. Moreover, if the firm bears the cost of the campaign, this option may not be profitable.

Our second set of counterfactuals considers instead revenue-neutral price changes that subsidize the price of the first surgery and tax the second. These interventions unequivocally lead to welfare improvements, which is of interest to both the firm's managers and policymakers.

Our findings suggest that uncertainty in these interactions is large and heterogeneous across patients, which in turn makes subsidizing initial take-up more efficient than implementing persuasive advertising. This suggests that even in a setting with limited repeat purchases, the value of revealed uncertainty may allow for price interventions that are profitable for the firm and increase consumer welfare.

Declarations

Funding and competing interests

The authors have no funding to report. This study was conducted in partnership with a healthcare firm in Mexico City that provided access to data and institutional background. The firm wishes to remain anonymous and had no involvement in the research design, analysis, or interpretation of results. They were not involved in drafting or reviewing the manuscript. The authors certify

that they have no financial or non-financial conflicts of interest related to this study.

References

Adhvaryu, Achyuta, Emilio Gutierrez, Anant Nyshadham, and Jorge Tamayo. 2020. "Diagnosing Quality: Learning, Amenities, and the Demand for Health Care." 1–34.

Arcidiacono, Peter, and Paul B. Ellickson. 2011. "Practical Methods for Estimation of Dynamic Discrete Choice Models." *Annual Review of Economics, Vol 3*, 3(2011): 363–394.

Aron-Dine, Aviva, Liran Einav, and Amy Finkelstein. 2013. "The RAND health insurance experiment, three decades later." *Journal of Economic Perspectives*, 27(1): 197–222.

Bergemann, Dirk, and Juuso Välimäki. 2006. "Dynamic pricing of new experience goods." *Journal of Political Economy*, 114(4): 713–743.

Berger, Jonah, Alan T. Sorensen, and Scott J. Rasmussen. 2010. "Positive effects of negative publicity: When negative reviews increase sales." *Marketing Science*, 29(5): 815–827.

Briesen, Sebastian, Helen Roberts, Dunera Ilako, Jefitha Karimurio, and Paul Courtright. 2010. "Are blind people more likely to accept free cataract surgery? A study of vision-related quality of life and visual acuity in Kenya." *Ophthalmic epidemiology*, 17(1): 41–49.

Busbee, Brandon G, Melissa M Brown, Gary C Brown, and Sanjay Sharma. 2003. "Cost-utility analysis of cataract surgery in the second eye." *Ophthalmology*, 110(12): 2310–2317.

Çakir, Hatice, Gülden Küçükakça Çelik, and Rabiye Çirpan. 2021. "Correlation between social support and psychological resilience levels in patients undergoing colorectal cancer surgery: a descriptive study." *Psychology, Health & Medicine*, 26(7): 899–910.

Castells, Xavier, Mercè Comas, Jordi Alonso, Mireia Espallargues, Vicente Martínez, Josep García-Arumí, and Miguel Castilla. 2006. "In a randomized controlled trial, cataract surgery in both eyes increased benefits compared to surgery in one eye only." *Journal of clinical epidemiology*, 59(2): 201–207.

Cataract surgical rates. 2017. *Community Eye Health*, 30(100): 88–89.

Cheung, D., and S. Sandramouli. 2005. "The consent and counselling of patients for cataract surgery: A prospective audit." *Eye*, 19(9): 963–971.

Ching, Andrew, and Masakazu Ishihara. 2010. "The effects of detailing on prescribing decisions under quality uncertainty." *QME*, 8: 123–165.

Ching, Andrew T. 2010. "Consumer learning and heterogeneity: Dynamics of demand for prescription drugs after patent expiration." *International Journal of Industrial Organization*, 28(6): 619–638.

Ching, Andrew T, Robert Clark, Ignatius Horstmann, and Hyunwoo Lim. 2016. "The effects of publicity on demand: The case of anti-cholesterol drugs." *Marketing Science*, 35(1): 158–181.

Ching, Andrew T, Tülin Erdem, and Michael P Keane. 2017. *Empirical models of learning dynamics: A survey of recent developments.* Springer.

Congdon, Nathan, and Ravi Thomas. 2014. "How can we solve the problem of low uptake of cataract surgery?" *Ophthalmic Epidemiology*, 21(3): 135–137.

Crawford, Gregory S, and Matthew Shum. 2005. "Uncertainty and learning in pharmaceutical demand." *Econometrica*, 73(4): 1137–1173.

Dardanoni, Valentino, and Adam Wagstaff. 1990. "Uncertainty and the demand for medical care." *Journal of Health Economics*, 9(1): 23–38.

Dickstein, Michael J. 2021. "Efficient provision of experience goods: Evidence from antidepressant choice." *NYU Stern Working Paper*.

Duarte, Fabian. 2012. "Price elasticity of expenditure across health care services." *Journal of health economics*, 31(6): 824–841.

Dupas, Pascaline. 2014. "Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment." *Econometrica*, 82(1): 197–228.

Dupas, Pascaline, and Edward Miguel. 2017. "Impacts and determinants of health levels in low-income countries." In *Handbook of economic field experiments*. Vol. 2, 3–93. Elsevier.

Ehrlich, Joshua R, Jacqueline Ramke, David Macleod, Helen Burn, Chan Ning Lee, Justine H Zhang, William Waldock, Bonnielin K Swenor, Iris Gordon, Nathan Congdon, et al. 2021. "Association between vision impairment and mortality: a systematic review and meta-analysis." *The Lancet Global Health*, 9(4): e418–e430.

ENSANUT. 2020. "Encuesta Nacional de Salud y Nutrición 2018-19: Resultados Nacionales." Instituto Nacional de Salud Pública, Cuernavaca, México.

Finkelstein, Amy, Petra Persson, Maria Polyakova, and Jesse M. Shapiro. 2021. "A Taste of Their Own Medicine: Guideline Adherence and Access to Expertise." *SSRN Electronic Journal*, 24(16): 63.

Foubert, Bram, and Els Gijsbrechts. 2016. "Try it, you'll like it—or will you? The perils of early free-trial promotions for high-tech service adoption." *Marketing Science*, 35(5): 810–826.

Gheini, Alireza, Amir Shakarami, Parsa Namdari, Mehrdad Namdari, and Ali Pooria. 2021. "Frequency of recurrence of peripheral artery disease among angioplasty and stenting patients." *Annals of Medicine and Surgery*, 72: 103146.

Gogate, P., S. Kulkarni, S. Krishnaiah, R. Deshpande, S. Joshi, A. Palimkar, and M. Deshpande. 2005. "Safety and Efficacy of Phacoemulsification Compared with Manual Small-Incision Cataract Surgery by a Randomized Controlled Clinical TrialSix-Week Results." *Ophthalmology*, 112(5): 869–874.

Gowrisankaran, Gautam, and Marc Rysman. 2012. "Dynamics of consumer demand for new durable goods." *Journal of Political Economy*, 120(6): 1173–1219.

Greene, William. 2004. "The behaviour of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects." *The Econometrics Journal*, 7(1): 98–119.

Gruber, Jonathan, and Maria Owings. 1996. "Physician Financial Incentives and Cesarean Section Delivery." *The RAND Journal of Economics*, 27(1): 99.

Gruener, Anna. 2015. "The effect of cataracts and cataract surgery on Claude Monet." *British Journal of General Practice*, 65(634): 254–255.

Hashemi, Hassan, Reza Pakzad, Abbasali Yekta, Mohamadreza Aghamirsalim, Mojgan Pakbin, Shahroukh Ramin, and Mehdi Khabazkhoob. 2020. "Global and regional prevalence of age-related cataract: a comprehensive systematic review and

meta-analysis." *Eye*, 34(8): 1357–1370.

Heemraz, B Sanjeev, Chan Ning Lee, Pirro G Hysi, Carole A Jones, Christopher J Hammond, and Omar A Mahroo. 2016. "Changes in quality of life shortly after routine cataract surgery." *Canadian Journal of Ophthalmology*, 51(4): 282–287.

Henderson, Bonnie An, and Julia Schneider. 2012. "Same-day cataract surgery should not be the standard of care for patients with bilateral visually significant cataract." *Survey of Ophthalmology*, 57(6): 580–583.

Hilger, James, Greg Rafert, and Sofia Villas-Boas. 2011. "Expert opinion and the demand for experience goods: An experimental approach in the retail wine market." *Review of Economics and Statistics*, 93(4): 1289–1296.

Javitt, Jonathan C, Earl P Steinberg, Phoebe Sharkey, Oliver D Schein, James M Tielsch, Marie Diener West, Marcia Legro, and Alfred Sommer. 1995. "Cataract surgery in one eye or both: a billion dollar per year issue." *Ophthalmology*, 102(11): 1583–1593.

Jing, Bing. 2011. "Social learning and dynamic pricing of durable goods." *Marketing Science*, 30(5): 851–865.

Johnson, Erin M., and M. Marit Rehavi. 2016. "Physicians treating physicians: Information and incentives in childbirth." *American Economic Journal: Economic Policy*, 8(1): 115–141.

Keel, Stuart, Andreas Müller, Sandra Block, Rupert Bourne, Matthew J Burton, Somnath Chatterji, Mingguang He, Van C Lansingh, Wanjiku Mathenge, Silvio Mariotti, et al. 2021. "Keeping an eye on eye care: monitoring progress towards effective coverage." *The Lancet Global Health*, 9(10): e1460–e1464.

Kowalski, Amanda. 2016. "Censored quantile instrumental variable estimates of the price elasticity of expenditure on medical care." *Journal of Business & Economic Statistics*, 34(1): 107–117.

Laidlaw, Alistair, and Richard Harrad. 1993. "Can second eye cataract extraction be justified?" *Eye*, 7(5): 680–686.

Laidlaw, DAH, RA Harrad, CD Hopper, A Whitaker, JL Donovan, ST Brookes, GW Marsh, TJ Peters, JM Sparrow, and SJ Frankel. 1998. "Randomised trial of effectiveness of second eye cataract surgery." *The Lancet*, 352(9132): 925–929.

Lansingh, Van C, Serge Resnikoff, Kimberly Tingley-Kelley, María E Nano, Marion Martens, Juan C Silva, Rainald Duerksen, and Marissa J Carter. 2010. "Cataract surgery rates in Latin America: a four-year longitudinal study of 19 countries." *Ophthalmic epidemiology*, 17(2): 75–81.

Lee, Ji Yea, Yeonsoo Jang, Sanghee Kim, and Woo Jin Hyung. 2020. "Uncertainty and unmet care needs before and after surgery in patients with gastric cancer: a survey study." *Nursing & Health Sciences*, 22(2): 427–435.

Lewallen, Susan, and Paul Courtright. 2000. "Recognising and reducing barriers to cataract surgery." *Community Eye Health*, 13(34): 20.

Lien, Chin-Yen, Hung-Ru Lin, Ing-Tiau Kuo, and Mei-Ling Chen. 2009. "Perceived uncertainty, social support and psychological adjustment in older patients with cancer being treated with surgery." *Journal of clinical nursing*, 18(16): 2311–2319.

Liu, Tianyu, Nathan Congdon, Xixi Yan, Ling Jin, Ying Wu, David Friedman, and Mingguang He. 2012. "A randomized, controlled trial of an intervention promoting cataract surgery acceptance in rural China: the Guangzhou Uptake of Surgery Trial

(GUSTO)." *Investigative ophthalmology & visual science*, 53(9): 5271–5278.

Lundström, Mats, Pik-Pin Goh, Ype Henry, Mohamad A Salowi, Peter Barry, Sonia Manning, Paul Rosen, and Ulf Stenevi. 2015. "The changing pattern of cataract surgery indications: a 5-year study of 2 cataract surgery databases." *Ophthalmology*, 122(1): 31–38.

Lurie, Nicole, Caren J Kamberg, Robert H Brook, Emmet B Keeler, and Joseph P Newhouse. 1989. "How free care improved vision in the health insurance experiment." *American journal of public health*, 79(5): 640–642.

Mailu, Eunice Wandia, Bhavisha Virendrakumar, Stevens Bechange, Emma Jolley, and Elena Schmidt. 2020. "Factors associated with the uptake of cataract surgery and interventions to improve uptake in low-and middle-income countries: A systematic review." *PLoS One*, 15(7): e0235699.

Manski, Charles F. 2018. "Reasonable patient care under uncertainty." *Health Economics*, 27(10): 1397–1421.

Maurer, Jürgen, and Katherine M Harris. 2016. "Learning to Trust Flu Shots: Quasi-Experimental Evidence from the 2009 Swine Flu Pandemic." *Health economics*, 25(9): 1148–1162.

Mercer, Gareth D, Penny Lyons, and Ken Bassett. 2019. "Interventions to improve gender equity in eye care in low-middle income countries: A systematic review." *Ophthalmic epidemiology*, 26(3): 189–199.

Miller, Kevin M., Thomas A. Oetting, James P. Tweeten, Kristin Carter, Bryan S. Lee, Shawn Lin, Afshan A. Nanji, Neal H. Shorstein, and David C. Musch. 2022. "Cataract in the Adult Eye Preferred Practice Pattern." *Ophthalmology*, 129(1): P1–P126.

Narayanan, Sridhar, and Puneet Manchanda. 2009. "Heterogeneous learning and the targeting of marketing communication for new products." *Marketing science*, 28(3): 424–441.

Nelson, Phillip. 1970. "Information and Consumer Behavior." *Journal of Political Economy*, 78(2): 311–329.

Newhouse, Joseph P. 1993. *Free for all?: lessons from the RAND health insurance experiment*. Harvard University Press.

OECD. 2016. "OECD Reviews of Health Systems: Mexico 2016." Paris:OECD Publishing.

Oster, Emily, and Rebecca Thornton. 2012. "Determinants of technology adoption: Peer effects in menstrual cup take-up." *Journal of the European Economic Association*, 10(6): 1263–1293.

Petrin, Amil, and Kenneth Train. 2010. "A Control Function Approach to Endogeneity in Consumer Choice Models." *Journal of Marketing Research*, XLVI(February): 3–13.

Reinstein, David A., and Christopher M. Snyder. 2005. "The influence of expert reviews on consumer demand for experience goods: A case study of movie critics." *Journal of Industrial Economics*, 53(1): 27–51.

Riaz, Yasmin, Samantha R de Silva, and Jennifer R. Evans. 2013. "Manual small incision cataract surgery (MSICS) with posterior chamber intraocular lens versus phacoemulsification with posterior chamber intraocular lens for age-related cataract." *Cochrane Database of Systematic Reviews*, 2014(11).

Shekhawat, Nakul S, Michael V Stock, Elizabeth F Baze, Mary K Daly, David E Vollman, Mary G Lawrence, and Amy S Chomsky. 2017. "Impact of first eye versus second eye cataract surgery on visual function and quality of life." *Ophthalmology*,

124(10): 1496–1503.

Sunada, Takeaki. 2020. "Consumer learning in a durable-goods environment and profitable free trials." 1–52.

Syed, Alishbah, Sarah Polack, Cristina Eusebio, Wanjiku Mathenge, Zakia Wadud, AKM Mamunur, Allen Foster, and Hannah Kuper. 2013. "Predictors of attendance and barriers to cataract surgery in Kenya, Bangladesh and the Philippines." *Disability and rehabilitation*, 35(19): 1660–1667.

Vuong, Quang H. 1989. "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses." *Econometrica*, 57(2): 307–333.

Wosinska, Marta. n.d.. "Just What the Patient Ordered? Direct-to-Consumer Advertising and the Demand for Pharmaceutical Products."

Yu, Man, Laurens Debo, and Roman Kapuscinski. 2016. "Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods." *Management Science*, 62(2): 410–435.

Zhang, Xiu Juan, Yuan Bo Liang, Ying Peng Liu, Vishal Jhanji, David C Musch, Yi Peng, Chong Ren Zheng, Hui Xi Zhang, Ping Chen, Xin Tang, et al. 2013. "Implementation of a free cataract surgery program in rural China: a community-based randomized interventional study." *Ophthalmology*, 120(2): 260–265.

Zhang, Yingzi. 2017. "Uncertainty in illness: theory review, application, and extension." *Number 6/November 2017*, 44(6): 645–649.

A Details on ocular health measures

TABLE 5: Description of ocular health measures

Measure	Description
Far visual acuity	A measure of a person's vision, typically set at "optical infinity", which is approximated at 20 feet.
Near visual acuity	A measure of a person's vision, defined as a comfortable reading distance of around 18 inches.
Amblyopia	Vision that does not develop properly during childhood. Also referred to as lazy eye.
Anisometropia	A condition in which the eyes have unequal refractive power, meaning the degree to which the lens converges or diverges light.
Astigmatism	Distorted shape of the cornea and/or lens that causes improper light refraction, leading to blurry and distorted vision of both near and far objects.
Myopia	Refractive error caused by the eye not focusing light properly on the retina, leading to blurry vision for distant objects. Also called nearsightedness.
Presbyopia	An increased rigidity in the lens caused by aging, which leads the eye to lose the ability to see things clearly up close.
Hypermetropia	A refractive error due to an eye focusing problem that causes close objects to appear blurred. Also called hyperopia or farsightedness.
Emmetropia	A state of vision without refractive error, leading to a sharp focus of objects that are far away, typically around 20 feet. The presence of a refractive error of this type is called ametropia.

Notes: This table describes the ocular health measures included in the summary statistics of main text Table 1. Information in this table is taken directly from the American Academy of Ophthalmology, www.aao.org.

B Alternative learning assumption for the simple model

Here, we make an alternative assumption of what "no learning" means. *With learning*, we assume that $\alpha_{i1} = \alpha_{i2} \equiv \alpha_i$, which is unknown before the first surgery, but known immediately afterwards. *Without learning*, we keep assuming that α_{it} are iid and are never known before the surgery.

Assume also that the firm is a monopolist with zero marginal costs that can charge consumers their willingness to pay, p_{it}, at each $t = 1, 2$. At $t = 2$, the consumer purchases if and only if $\mathbb{E}_\alpha [u_{i2} - p_{i2}] \geq 0$, which implies that $p_{i2} = \max\{0, \mathbb{E}_\alpha [u_{i2}]\}$, and depends on learning thus:

$$p_{i2} = \begin{cases} \max\{0, \alpha_i + \varepsilon_{i2}\} & \text{with learning} \\ \max\{0, \varepsilon_{i2}\} & \text{without learning} \end{cases}$$

because, with learning, α_{i2} is already known at $t = 2$. At $t = 1$ the consumer purchases if and only if

$$\mathbb{E}_\alpha [u_{i1} - p_{i1}] + \mathbb{E}_{\alpha,\varepsilon} [u_{i2} - p_{i2} | u_{i2} - p_{i2} > 0] > 0,$$

where the second term represents an option value at $t = 1$, which is zero in equilibrium. Analogously, the price at $t = 1$ becomes

$$p_{i1} = \max\{0, \varepsilon_{i1}\} \qquad \text{with or without learning,}$$

because, with either learning assumption, at $t = 1$ the consumer has an option value, but doesn't know about α. All prices are truncated random variables. Under normality, it can be shown that, as a function of σ_α^2,

$$\frac{\mathbb{V}[p_{i2}]}{\mathbb{V}[p_{i1}]} = \begin{cases} 1 + \sigma_\alpha^2 & \text{with learning} \\ 1 & \text{without learning.} \end{cases}$$

C Simple model with risk preferences and diminishing marginal utility

Here, we explicitly allow for risk aversion with a simple modification of the utility function, which exhibits constant absolute risk aversion. In this case, an

individual i receives a utility $u_{it} = \alpha_i + 1 - e^{\rho\alpha_i} - m_{it} + \varepsilon_{it}$ from undergoing surgery $t = 1, 2$, where $m_{i1} = 0$ and $m_{i2} = m$ represent a diminishing marginal utility from the second surgery, and ρ is a risk preference, with $\rho > 0$ to represent risk aversion, risk neutrality corresponds to $\rho = 0$, and risk seeking is $\rho < 0$.

At $t = 2$, the consumer purchases if and only if $u_{i2} - p_{i2} > 0$, which implies that

$$p_{i2} = \max\{0, \alpha_i + 1 - e^{\rho\alpha_i} - m + \varepsilon_{i2}\},$$

with both full information or learning.

At $t = 1$ the consumer purchases if and only if

$$\mathbb{E}_\alpha \left[u_{i1} - p_{i1}\right] + \mathbb{E}_{\alpha,\varepsilon} \left[u_{i2} - p_{i2} | u_{i2} - p_{i2} \geq 0\right] P \left[u_{i2} - p_{i2} \geq 0\right] \geq 0,$$

where the second term represents an option value at $t = 1$, which is present regardless of the informational assumption, but in the equilibrium of this simple model is equal to 0.

Therefore, the price at $t = 1$ is

$$p_{i1} = \begin{cases} \max\{0, \mathbb{E}_\alpha \left[\alpha_i + 1 - e^{\rho\alpha_i}\right] + \varepsilon_{i1}\} & \text{with learning,} \\ \max\{0, \alpha_i + 1 - e^{\rho\alpha_i} + \varepsilon_{i1}\} & \text{with full information,} \end{cases}$$

because, with full information, the consumer knows α_i at $t = 1$.

Lemma 2. Suppose that $\varepsilon \sim \mathcal{N}(0, 1)$, $\alpha \sim \mathcal{N}(0, \sigma_\alpha)$, and are independent. Let $m_{i1} = 0$, $m_{i2} = m > 0$ and $\rho \in \mathbb{R}$. Then, the ratio

$$\frac{\mathbb{V}\left[p_{i2}\right]}{\mathbb{V}\left[p_{i1}\right]} = \begin{cases} \text{is increasing in } \sigma_\alpha & \text{with learning} \\ \text{is constant in } \sigma_\alpha & \text{with full information.} \end{cases}$$

And, in particular, $\frac{\mathbb{V}[p_{i2}]}{\mathbb{V}[p_{i1}]} > 1$ with learning.

The proof follows from the fact that at $m = 0$, price variances are equal with full information, but, with learning, the variances of p_{i1} is that of a truncated standard normal. For $m > 0$, m is a constant, which decreases the variance of p_{i2}, but not as a function of σ_α. In other words, the ratio of variances is roughly constant with full information, but increasing in σ_α with learning. Figure 9 shows numerical simulations. Note that the results do not depend on $\rho > 0$. We still have the same qualitative results for a risk seeking population.

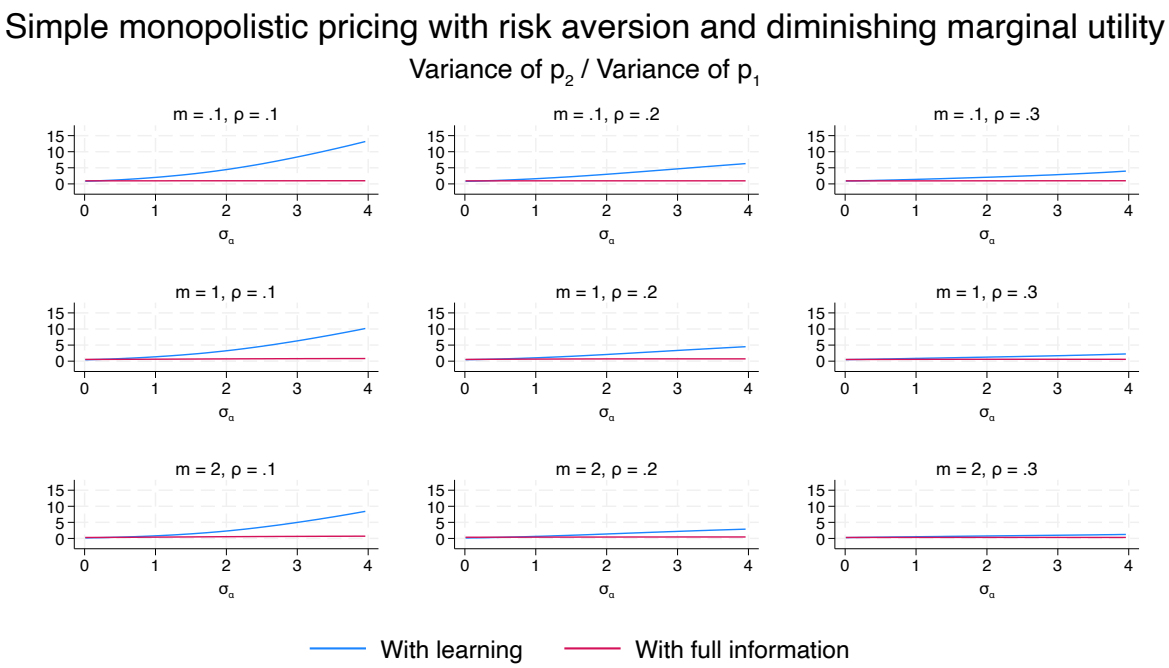


FIGURE 9: Predicted behavior of price dispersion standard deviations

Note: In this exercise, $\varepsilon \sim \mathcal{N}(0,1)$ and $\alpha \sim \mathcal{N}(0,\sigma_\alpha)$. Utility of second period penalized by m. Risk aversion parameter is ρ. Equilibrium prices are described in section C.

D A simple model without perfect price discrimination

Consider the same setup as in section 3.1.3, but, in this case, the monopolist does not know the realizations of ε_{it} and α_i, but knows their distributions and can price accordingly. The firm can still offer different prices to different people. In period $t = 2$, the agent receives $u_{i2} - p_{i2} = \alpha_i - \varepsilon_{i2} - p_{i2}$ where $\varepsilon \sim \mathcal{N}(0,1)$ and $\alpha \sim \mathcal{N}(1,\sigma_\alpha)$. The firm maximizes profits by choosing p_{i2}: $\max_{p_{i2}} p_{i2} P\left[u_{i2} - p_{i2} > 0\right]$, which implies the optimal price solves

$$\frac{p_{i2}}{\sqrt{1+\sigma_\alpha^2}} = \frac{1 - \Phi\left(\frac{p_{i2}}{\sqrt{1+\sigma_\alpha^2}}\right)}{\phi\left(\frac{p_{i2}}{\sqrt{1+\sigma_\alpha^2}}\right)},$$

which can be shown to be increasing in σ_α.

At $t = 1$ the consumer purchases if and only if

$$\mathbb{E}_\alpha \left[u_{i1} - p_{i1} \right] + \mathbb{E}_{\alpha,\varepsilon} \left[u_{i2} - p_{i2} | u_{i2} - p_{i2} \geq 0 \right] P \left[u_{i2} - p_{i2} \geq 0 \right] \geq 0,$$

where the second term represents an option value at $t = 1$, which depends on the informational assumption. Let v_i represent this option value. With full information, the option value depends on α: $v_i(\alpha_i)$. However, with learning, the option value v_i is constant in α_i. In any case, the firm's problem is $\max_{p_{i1}} p_{i1} P \left[\mathbb{E}_\alpha \left[\alpha_i \right] + \varepsilon_{i1} + v_i - p_{i1} > 0 \right]$, where the expectation depends on the informational assumption.

With full information, we can solve this program numerically, or we can approximate $v_i(\alpha_i) \approx v_i(0) + v_i'(0)\alpha_i$ in order to compute the optimal pricing. With this approximation, the optimal p_{i1} solves

$$\frac{p_{i1}}{\sqrt{1 + \sigma_\alpha^2 (1 + v_i'(0))^2}} = \frac{1 - \Phi\left(\frac{p_{i1} - v_i(0)}{\sqrt{1 + \sigma_\alpha^2 (1 + v_i'(0))^2}} \right)}{\phi\left(\frac{p_{i1} - v_i(0)}{\sqrt{1 + \sigma_\alpha^2 (1 + v_i'(0))^2}} \right)},$$

where $v_i(\alpha) = (\alpha - p_{i2} + \lambda(\alpha - p_{i2}))\Phi(\alpha - p_{i2})$. With full information, this price can be shown to be also increasing in σ_α.

On the other hand, with learning, the optimal pricing solves

$$p_{i1} = \frac{1 - \Phi\left(p_{i1} - v_i \right)}{\phi\left(p_{i1} - v_i \right)},$$

where $v_i \equiv (-p_{i2} + \sqrt{1 + \sigma^2}\lambda(-p_{i2}/\sqrt{1 + \sigma^2}))\Phi(-p_{i2}/\sqrt{1 + \sigma^2})$ can be shown to be increasing in σ_α, but at a lower rate.

Figure 10 shows a numerical simulation where we plot the optimal price ratios. Under full information, the price ratios are roughly constant and equal to 1, but with learning, the price of the second surgery increases relative to the first one when σ_α grows, in line with section 3.1.3.

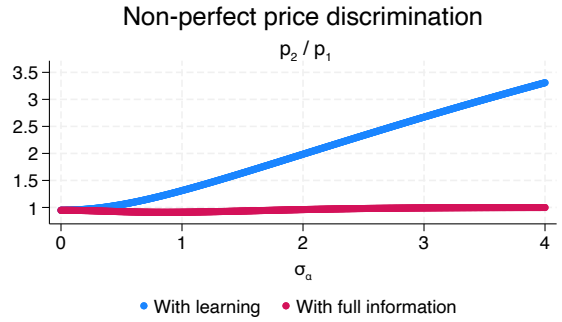
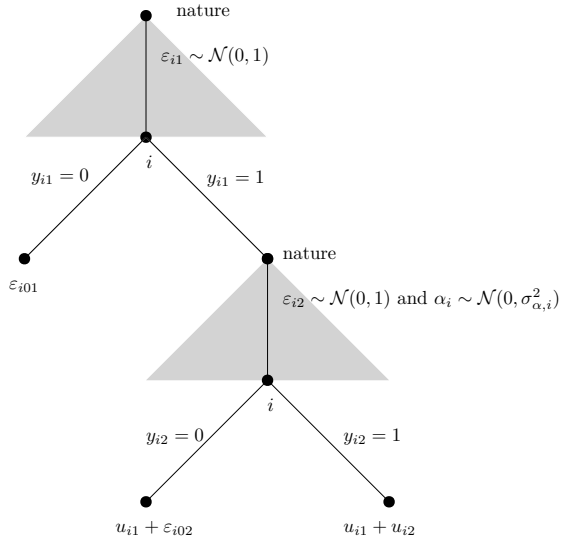


FIGURE 10: Predicted behavior of prices

Note: In this exercise, $\varepsilon \sim \mathcal{N}(0,1)$ and $\alpha \sim \mathcal{N}(0,\sigma_\alpha)$. Equilibrium prices are described in section D.



1. Nature draws shock ε_{i1}
2. Given ε_{i1}, i operates or not
3. Nature draws shocks $\varepsilon_{i2}, \alpha_i$
4. Given $\varepsilon_{i2}, \alpha_i$, i operates or not
5. Payoffs

$$u_{i1} = \alpha_i + \boldsymbol{\beta}' \boldsymbol{x}_{i1} + \varepsilon_{i1}$$
$$u_{i2} = \alpha_i + \boldsymbol{\beta}' \boldsymbol{x}_{i2} + \varepsilon_{i2}$$

FIGURE 11: Model tree and basic timing

E Model tree and demand derivations

Assuming normal errors (assumption 1), we can simplify,

$$\begin{aligned}
&\mathbb{E}\left[u_{i2} - u_{i02} | u_{i2} - u_{i02} > 0\right] \\
&= \boldsymbol{\beta}' \boldsymbol{x}_{i2} + \mu_{\alpha,i} + \sqrt{1+\sigma^2_{\alpha,i}} \mathbb{E}\left[\frac{\alpha_i + \varepsilon_{i2} - \varepsilon_{i02} - \mu_{\alpha,i}}{\sqrt{1+\sigma^2_{\alpha,i}}} \middle| \frac{\alpha_i + \varepsilon_{i2} - \varepsilon_{i02} - \mu_{\alpha,i}}{\sqrt{1+\sigma^2_{\alpha,i}}} > -\frac{\boldsymbol{\beta}' \boldsymbol{x}_{i2} + \mu_{\alpha,i}}{\sqrt{1+\sigma^2_{\alpha,i}}} \right], \\
&= \boldsymbol{\beta}' \boldsymbol{x}_{i2} + \mu_{\alpha,i} + \sqrt{1+\sigma^2_{\alpha,i}} \lambda \left(\frac{\boldsymbol{\beta}' \boldsymbol{x}_{i2} + \mu_{\alpha,i}}{\sqrt{1+\sigma^2_{\alpha,i}}} \right), \qquad (4)
\end{aligned}$$

conditional on the first surgery. Also,

$$\begin{aligned} P\left[u_{i2}-u_{i02}>0\right] &= P\left[\alpha_i+\varepsilon_{i2}-\varepsilon_{i02}+\boldsymbol{\beta}'\boldsymbol{x}_{i2}>0\right] \\ &= \Phi\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2}+\mu_{\alpha,i}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right). \end{aligned}$$

Then,

$$\begin{aligned} &\mathbb{E}_{\alpha_i}[u_{i1}-u_{i01}+\mathbb{E}\left[u_{i2}-u_{i02}\right]] \\ &= \mu_{\alpha,i}+\boldsymbol{\beta}'\boldsymbol{x}_{i1}+\varepsilon_{i1}-\varepsilon_{i01} \\ &+\left\{\boldsymbol{\beta}'\boldsymbol{x}_{i2}+\mu_{\alpha,i}+\sqrt{1+\sigma_{\alpha,i}^2}\lambda\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2}+\mu_{\alpha,i}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right)\right\}\Phi\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2}+\mu_{\alpha,i}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right), \end{aligned}$$

and,

$$\begin{aligned} &P\left[y_{i1}=1\right] \\ &= P\left[\mathbb{E}_{\alpha_i}[u_{i1}-u_{i01}+\mathbb{E}\left[u_{i2}-u_{i02}\right]]>0\right] \\ &= \Phi\left[\mu_{\alpha,i}+\boldsymbol{\beta}'\boldsymbol{x}_{i1}+\left\{\boldsymbol{\beta}'\boldsymbol{x}_{i2}+\mu_{\alpha,i}+\sqrt{1+\sigma_{\alpha,i}^2}\lambda\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2}+\mu_{\alpha,i}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right)\right\}\Phi\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2}+\mu_{\alpha,i}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right)\right]. \end{aligned}$$

Also, once α_i is known to the consumer, the unconditional demand for the second surgery is

$$\begin{aligned} P\left[y_{i2}=1\right] &= P\left[y_{i2}=1,y_{i1}=1\right]+P\left[y_{i2}=1,y_{i1}=0\right], \\ &= P\left[y_{i2}=1|y_{i1}=1\right]P\left[y_{i1}=1\right]+0, \end{aligned}$$

and

$$P\left[y_{i2}=1|y_{i1}=1\right] = P\left[u_{i2}-u_{i02}>0\right] = \Phi\left(\frac{\mu_{\alpha,i}+\boldsymbol{\beta}'\boldsymbol{x}_{i2}}{\sqrt{1+\sigma_{\alpha,i}^2}}\right)$$

The probability mass function of both operations (y_{i1},y_{i2}) is

$$P\left[y_{i1}=y,y_{i2}=y'\right] = P\left[y_{i2}=y'|y_{i1}=y\right]P\left[y_{i1}=y\right],$$

where $y,y'=0$ or 1.

Then, for each i the likelihood of observing (y_{i1}, y_{i2}) is

$$
\begin{aligned}
P\left[y_{i1}=0, y_{i2}=1\right] &= P\left[y_{i2}=1 | y_{i1}=0\right] P\left[y_{i1}=0\right] = 0 \\
P\left[y_{i1}=0, y_{i2}=0\right] &= P\left[y_{i2}=0 | y_{i1}=0\right] P\left[y_{i1}=0\right] = 1 - s_{i1} \\
P\left[y_{i1}=1, y_{i2}=0\right] &= P\left[y_{i2}=0 | y_{i1}=1\right] P\left[y_{i1}=1\right] = s_{i1}(1 - s_{i2}) \\
P\left[y_{i1}=1, y_{i2}=1\right] &= P\left[y_{i2}=1 | y_{i1}=1\right] P\left[y_{i1}=1\right] = s_{i1} s_{i2}.
\end{aligned}
$$

Or, equivalently,

$$
s_{i1}^{y_{i1}}(1 - s_{i1})^{1-y_{i1}} \left[s_{i2}^{y_{i2}}(1 - s_{i2})^{1-y_{i2}} \right]^{y_{i1}}.
$$

Finally, because the expected demand of consumer i is $D_i = s_{i1} + s_{i1}s_{i2}$, price elasticities of demand are defined as

$$
\epsilon_{i1} \equiv \frac{\partial s_{i1}}{\partial p_{i1}} \frac{p_{i1}}{s_{i1}} \quad \text{and} \quad \epsilon_{i2} \equiv \frac{\partial s_{i1}s_{i2}}{\partial p_{i2}} \frac{p_{i2}}{s_{i1}s_{i2}}.
$$

It can be shown that

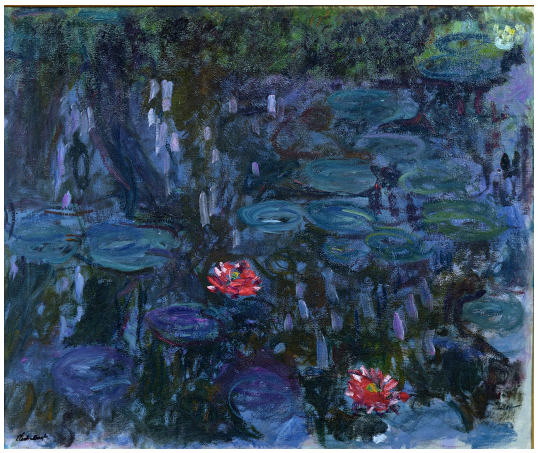
$$
\frac{\partial D_i}{\partial p_{i1}} \frac{p_{i1}}{D_i} = \epsilon_{i1} \quad \text{and} \quad \epsilon_{i2} = \frac{\partial s_{i2}}{\partial p_{i2}} \frac{p_{i2}}{s_{i2}} + \epsilon_{i1} s_{i2}.
$$

F Claude Monet's α

Claude Monet, one of the most celebrated impressionist artists, developed cataracts that profoundly affected his vision and, in turn, his art. Gruener (2015) recounts Monet's experience, which we further use as a case study motivating our study. According to Gruener (2015), Monet's cataracts are evident when comparing works like "Waterlilies and Japanese Bridge" from the late 19th century to "Nymphéas Reflets de Saule" from around 1916, where his declining eyesight is evidenced in broader brush strokes and darker tones. By 1916, Monet's cataracts had significantly progressed, and though urged by friends and surgeons to undergo surgery, Monet was initially too fearful, consistent with how we model information prior to any surgery (though admittedly, surgeries were not as successful at the time). In 1923, Monet finally got surgery, although "his stubbornness" made recovery a challenge, also consistent with how we model expectations about post-operative care. The effect of cataracts on Monet's late work, shaped purely by his declining vision and his reluctance to get treatment, is seen by art historians as an unexpected bridge between impressionism and modern abstract art.



(A) Water lilies and Japanese bridge



(B) Nymphéas reflets de saule

FIGURE 12: Two of Monet's paintings: Giverny period c.1897 (left) and with cataracts c.1916 (right)

G Price endogeneity and control function

To deal with endogenous prices, we use a control function (Petrin and Train, 2010). We assume:

Assumption 3. Shocks can be decomposed as $\varepsilon - \varepsilon_0 = \gamma\rho + \widetilde{\varepsilon}$, where prices $p \perp \widetilde{\varepsilon}$, and ρ is correlated with prices, with $\mathbb{V}[\rho] = 1$.

Then,

$$\mathbb{V}[\varepsilon - \varepsilon_0] = 1 = \gamma^2 + \mathbb{V}[\widetilde{\varepsilon}] \Rightarrow \mathbb{V}[\widetilde{\varepsilon}] = 1 - \gamma^2.$$

Define

$$\sigma_{\widetilde{\varepsilon}} \equiv \sqrt{1 - \gamma^2}.$$

Therefore, by decomposing $\varepsilon - \varepsilon_0$ in the preceding derivations, we have

$$\begin{aligned} P[y_{i2} = 1] &= P\left[\alpha_i + \boldsymbol{\beta}'\boldsymbol{x}_{i2} + \gamma\rho_{i2} + \widetilde{\varepsilon}_{i2} > 0 | y_{i1} = 1\right] \\ &= \Phi\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2} + \gamma\rho_{i2}}{\sqrt{\sigma^2_{\widetilde{\varepsilon}} + \sigma^2_{\alpha,i}}}\right) \\ &= \Phi\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2} + \gamma\rho_{i2}}{\sigma_{\widetilde{\varepsilon}}\sqrt{1 + \frac{\sigma^2_{\alpha,i}}{\sigma^2_{\widetilde{\varepsilon}}}}}\right), \end{aligned}$$

and

$$\begin{aligned}
&\mathbb{E}\left[u_{i2} - u_{i02} | u_{i2} - u_{i02} > 0\right] \\
&= \boldsymbol{\beta}'\boldsymbol{x}_{i2} + \gamma\rho_{i2} + \mathbb{E}\left[\alpha_i + \widetilde{\varepsilon}_{i2} | \alpha_i + \boldsymbol{\beta}'\boldsymbol{x}_{i2} + \gamma\rho_{i2} + \widetilde{\varepsilon}_{i2} > 0\right] \\
&= \boldsymbol{\beta}'\boldsymbol{x}_{i2} + \gamma\rho_{i2} + \sigma_{\widetilde{\varepsilon}}\sqrt{1 + \frac{\sigma^2_{\alpha,i}}{\sigma^2_{\widetilde{\varepsilon}}}}\lambda\left(\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i2} + \gamma\rho_{i2}}{\sigma_{\widetilde{\varepsilon}}\sqrt{1 + \frac{\sigma^2_{\alpha,i}}{\sigma^2_{\widetilde{\varepsilon}}}}}\right).
\end{aligned}$$

Then,

$$\begin{aligned}
P\left[y_{i1} = 1\right] &= P\left[\mathbb{E}_{\alpha_i}\left[u_{i1} + \mathbb{E}\left[u_{i2} - u_{i02} | u_{i2} - u_{i02} > 0\right] P\left[y_{i2} = 1\right]\right] > 0\right] \\
&= P\left[\boldsymbol{\beta}'\boldsymbol{x}_{i1} + \gamma\rho_{i1} + \widetilde{\varepsilon}_{i1} + \mathbb{E}\left[u_{i2} - u_{i02} | u_{i2} - u_{i02} > 0\right] P\left[y_{i2} = 1\right] > 0\right], \\
&= \Phi\left[\frac{\boldsymbol{\beta}'\boldsymbol{x}_{i1} + \gamma\rho_{i1} + \mathbb{E}\left[u_{i2} - u_{i02} | u_{i2} - u_{i02} > 0\right] P\left[y_{i2} = 1\right]}{\sigma_{\widetilde{\varepsilon}}}\right].
\end{aligned}$$

Therefore, every parameter of the model is rescaled by $1/\sigma_{\widetilde{\varepsilon}}$, which needs to be accounted for to report the parameters in the original scale. Indeed, from an estimate of $\widehat{(\frac{\gamma}{\sigma_{\widetilde{\varepsilon}}})}$, we can back out

$$\widehat{\gamma} = \frac{\widehat{(\frac{\gamma}{\sigma_{\widetilde{\varepsilon}}})}}{\sqrt{1 + \widehat{(\frac{\gamma}{\sigma_{\widetilde{\varepsilon}}})}^2}} \Rightarrow \widehat{\sigma_{\widetilde{\varepsilon}}} = \sqrt{\frac{1}{1 + \widehat{(\frac{\gamma}{\sigma_{\widetilde{\varepsilon}}})}^2}}.$$

H First-stage regressions

We present the results from our first stage below. We obtain an F-statistic higher than 29, indicating that our instruments are relevant.

I Details on unobserved price predictions

In order to predict the missing price quotes on second surgeries, we employ a least absolute shrinkage selector operator (LASSO). Specifically, we predict log prices using:

- Patient's characteristics: age, gender, access to private insurance, social security, cataract scores, and ocular conditions, namely, ampliopia, anisometropia, astigmatism, myopia, presbyopia, hypermetropia, and emmetropia.

TABLE 6: FIRST STAGE OF DEMAND ESTIMATIONS

DEP VAR: log price_{it}	(1)
Sales agent's prior 15-day-avg (log) price	0.139
	(0.0254)
log Age	0.014
	(0.0217)
Female	−0.007
	(0.0066)
Min cataract score	0.009
	(0.0024)
Max cataract score	−0.006
	(0.0026)
Risk aversion proxy	−0.050
	(0.0043)
OTHER CONTROLS	YES
R_a^2	0.73
FIRST-STAGE IV'S F	29.91
PATIENTS	3,822
QUOTES	7,627

Notes: Standard errors clustered at individual level. Risk aversion proxy is (log) number of visits to obtain a price quote. Other controls include: sales agents fixed effects, surgery characteristics, and type of insurance as proxy for income.

- Surgery's characteristics: type of intraocular lens and type of surgery.
- Personnel: identity of sales agents, optometrists, and ophthalmologists who interacted with the patient.

These covariates amount to a total of 281 predictors, of which 133 were selected by LASSO. The penalty parameter was selected by cross-validation, using 10 folds. We use all observed price quotes for this estimation. We find a mean prediction error of .07, which is small, given the average log price is about 9.4.

Figure 13 shows the price histograms of observed prices for first and second surgeries, and the predicted price distribution for the unobserved second surgery prices. The graph and the mean prediction error give us confidence in our procedure to predict the missing data.

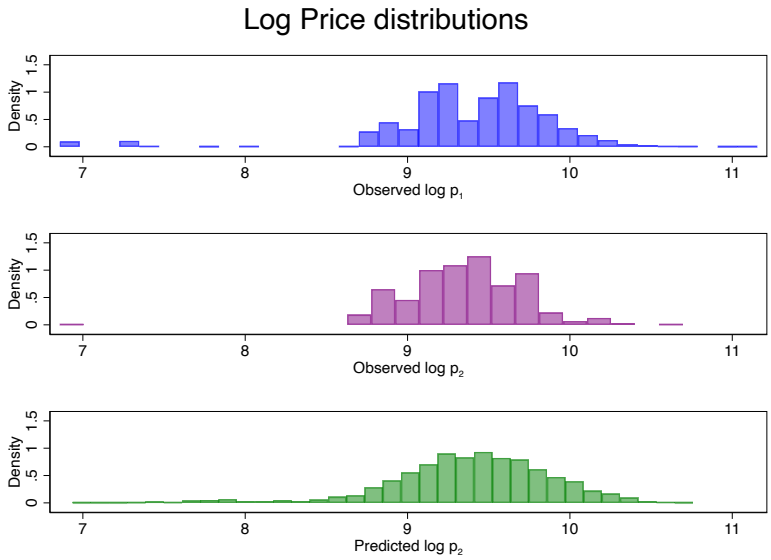


FIGURE 13: Observed and predicted (log) price distributions

J Model fit in a hold-out sample

We perform our estimations on a training set, consisting of a random sample of 50% of the data, and we test our estimations in the hold-out sample of the remaining 50%. We find our preferred specification outperforms the rest, and is not simply overfitting the data.

TABLE 7: DEMAND ESTIMATIONS IN TRAINING SET

DEP VAR: Operates$_{it}$	(1)	(2)	(3)
log price	−1.20	−2.22	−2.38
	(0.594)	(1.114)	(1.074)
log Age	0.03	−0.07	−0.12
	(0.085)	(0.144)	(0.246)
Female	−0.06	−0.14	−0.08
	(0.041)	(0.061)	(0.221)
Min cataract score	0.04	0.12	0.15
	(0.013)	(0.024)	(0.065)
Max cataract score	0.08	0.03	0.00
	(0.017)	(0.025)	(0.073)
Risk aversion proxy	0.74	1.08	1.03
	(0.061)	(0.227)	(0.213)
DEP VAR: $\sigma_{\alpha,i}$			
log Age			−0.05
			(0.269)
Female			−0.16
			(0.257)
Min cataract score			−0.12
			(0.142)
Max cataract score			−0.05
			(0.144)
RE health score			−0.03
			(0.144)
LE health score			−0.03
			(0.141)
cons		3.97	3.64
		(0.222)	(1.153)
ELASTICITIES			
ALL OPS	-1.28	-4.27	-4.03
FIRST OPS	-1.29	-5.94	-5.56
SECOND OPS	-1.28	-2.56	-2.45
OTHER CONTROLS	YES	YES	YES
CONTROLS ($\sigma_{\alpha,i}$)	NO	NO	YES
CONTROL FUNCTION	YES	YES	YES
MPE (TEST SET)	0.40	0.38	0.37
R^2_p (TEST SET)	0.03	0.07	0.10
FIRST-STAGE IV'S F	44.38	16.35	16.35
PATIENTS	1,873	1,873	1,873
QUOTES	3,739	3,739	3,739

Notes: Bootstrapped standard errors clustered at individual level with 100 repetitions. Models trained on a random subsample of 50% of the original data, and tested on the remaining 50%. Min and max cataract scores are a patient's eyes minimum and maximum score in a 0 to 6 scale. Control functions for prices are constructed with the 15-day rolling-average of log prices up to the moment as an instrument (Petrin and Train, 2010). Other controls include: sales agents fixed effects, surgery characteristics, and type of insurance as proxy for income. In the $\sigma_{\alpha,i}$ equation, eye health scores count the number of comorbidities (ampliopia, anisometropia, astigmatism, myopia, presbyopia, hypermetropia, and emmetropia) present in each eye. MPE stands for mean prediction error, computed on the test set. R^2_p stands for pseudo-R^2, constructed as described in footnote 27. Price elasticities with respect to unconditional demands: $P[y_{it} = 1]$, $t = 1, 2$.