DISCUSSION PAPER SERIES

DP16706

Price Discrimination and Big Data: Evidence from a Mobile Puzzle Game

Louis Pape, Christian Helmers, Alessandro Iaria, Stefan Wagner and Julian Runge

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

Price Discrimination and Big Data: Evidence from a Mobile Puzzle Game

Louis Pape, Christian Helmers, Alessandro Iaria, Stefan Wagner and Julian Runge

Discussion Paper DP16706
Published 09 November 2021
Submitted 28 October 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Price Discrimination and Big Data: Evidence from a Mobile Puzzle Game

Abstract

We use a unique dataset from a mobile puzzle game to investigate the welfare consequences of price discrimination. We rely on experimental variation to characterize player behavior and estimate a model of demand for game content. Our counterfactual simulations show that optimal uniform pricing would increase profit by +340% with respect to the game developer's observed pricing. This is almost the same as the increase in profit associated with first-degree price discrimination (+347%). All pricing strategies considered—including optimal uniform pricing—would induce a transfer of surplus from players to game developer without, however, generating sizeable dead-weight losses.

JEL Classification: D40, L11

Keywords: price discrimination, personalized pricing, mobile apps, online games, freemium

Louis Pape - Louis.PAPE@ensae.fr
CREST

Christian Helmers - christian.r.helmers@gmail.com
Santa Clara University

Alessandro Iaria - iaria.alessandro@gmail.com
University of Bristol and CEPR

Stefan Wagner - wagner.stefan@gmail.com
ESMT Berlin

Julian Runge - julian.runge@google.com
RWTH Aachen

Price Discrimination and Big Data: Evidence from a Mobile Puzzle Game*

Louis-Daniel Pape
CREST – Ecole Polytechnique

Christian Helmers
Santa Clara University

Alessandro Iaria
University of Bristol

Stefan Wagner
ESMT Berlin

Julian Runge
RWTH Aachen

November 9, 2021

ABSTRACT

We use a unique dataset from a mobile puzzle game to investigate the welfare consequences of price discrimination. We rely on experimental variation to characterize player behavior and estimate a model of demand for game content. Our counterfactual simulations show that optimal uniform pricing would increase profit by +340% with respect to the game developer's observed pricing. This is almost the same as the increase in profit associated with first-degree price discrimination (+347%). All pricing strategies considered—including optimal uniform pricing—would induce a transfer of surplus from players to game developer without, however, generating sizeable dead-weight losses.

KEYWORDS: Price discrimination, personalized pricing, mobile apps, online games, freemium.

JEL Classification: D40, L11

*We acknowledge helpful comments from seminar participants at CREST, ESMT, Telecom ParisTech, and the University of Linz. We are grateful to Lea Bignon for excellent research assistance. We are particularly grateful to Manasa Patnam for her invaluable input and advice.

1 Introduction

Price discrimination is ubiquitous in offline markets (Varian, 1989). This is not surprising since charging different prices for homogeneous goods to different consumers according to their willingness-to-pay enables companies to increase profits (Tirole, 1988).

Surprisingly, however, so far price discrimination has not been implemented on a large scale in the monetization of digital products, where "examples remain fairly limited" (Council of Economic Advisors, 2015). This is surprising because price discrimination could be implemented at comparably lower costs for digital products. By the nature of the underlying technology, companies can easily collect large amounts of detailed data about consumer characteristics and behavior (Goldfarb and Tucker, 2019). The resulting (big) data could be used to determine personalized prices for consumers with different characteristics and search or consumption histories, that is, engage in price discrimination, to increase a company's profit. That being said, firms have been reported to fear consumer backlash and negative press coverage when engaging in price discrimination (Garbarino and Maxwell, 2010; Li and Jain, 2016; DellaVigna and Gentzkow, 2019). An example is Amazon's early attempt to price discriminate buyers of DVDs based on their individual purchase history, which met dramatic resistance and negative publicity (Rosencrance, 2019). Similarly, the Wall Street Journal revealed in 2012 that Staples was charging consumers online different prices based on their location. However, there is a growing body of evidence to suggest that, in many important industries, because of path-dependence, imperfect information, learning, or conflicting incentives, sometimes for-profit firms do not maximize profit (Cho and Rust, 2010; DellaVigna and Gentzkow, 2019; Fioretti, 2020; Hortaçsu et al., 2021; Huang et al., 2020; Orbach and Einav, 2007) and, indeed, Dube and Misra (2019) estimate that, in the context of a digital recruiting firm, both profit *and* consumer surplus would often increase with personalized pricing.

In this paper, we contribute to this debate by investigating the welfare consequences of price discrimination for a mobile gaming app. Gaming accounts for over three quarters of total app revenue in the major app stores, including Apple's App Store and Google's Playstore (TechCrunch, June 11, 2019). Mobile games attracted over 150 million users worldwide and generated almost US$100 billion in revenue in 2020 (Statista, 2021). While game developers in this market collect around 90% of their revenues through paying customers on the basis of freemium models[1] (advertising accounts only for a small share of revenue), the efficiency of alternative pricing strategies in this industry has not yet received much scholarly attention. In our study, we focus on the popular category of "casual games" (games characterized by a sequence of levels that can be solved in a short amount of time) and analyze data from

[1]Freemium refers to a hybrid pricing model combining free and paid features of a product—basic features of a product can be used for free perpetually while more advanced features or more intensive use requires the payment of a fee. Freemium is particularly common for mobile apps, where consumers strongly favor apps that are free and monetized through in-app purchases rather than advertisement (Ghose and Han, 2014a).

one of the most popular match-3 games of all times. The data consist of the full in-game behavior and purchase decisions of about 300,000 players around the world for a two-week period between the end of October and early November 2013.

In the game, players solve puzzles through levels, and while the first 40 levels are free, from level 40 players must unlock a "pay-gate" every 20 levels to proceed (to the premium levels). Players can unlock any of these pay-gates by purchasing a "key." In this paper, we study the welfare consequences—for both the game developer and players—of five alternative pricing strategies to unlock pay-gates: the game developer's observed pricing; optimal uniform pricing; two forms of personalized pricing (third-degree price discrimination), one based on a player's gaming ability as measured in the free levels prior to the first pay-gate, the other based on the GDP per capita of a player's country; and first-degree price discrimination.

We combine experimental variation in the data with a structural model to estimate demand for free and paid-for content (additional levels) and then simulate the above counterfactual pricing strategies. We rely on the experimental variation in the data in two ways. First, we use it to learn about player behavior and document that players are unsophisticated and myopic, which greatly simplifies the specification of the structural model and the subsequent counterfactual simulations. Second, we rely on the experimental variation to estimate the demand model, in particular to address the standard challenges of price endogeneity and endogenous sample selection that would otherwise complicate identification (Gandhi and Nevo, 2021). Another helpful feature of the game is that no advertisement was displayed to players around the time of data collection. This allows us to focus on in-app purchases as the only source of revenue in the game. The co-existence of in-app purchases and advertisement would introduce dynamic interactions between pricing and advertising decisions which would be extremely hard to model, estimate, and ultimately simulate in counterfactual scenarios (Dubois et al., 2017).

Our counterfactual simulation results confirm that price discrimination will in general enable the game developer (a monopolist in our model) to seize larger portions of consumer surplus and thus increase profit (Varian, 1989). By comparing the relative performance of the counterfactual pricing strategies, we obtain three perhaps less obvious insights. First, observed pricing is far from optimal and the game developer could increase profit by +340% by limiting itself to the use of uniform pricing. A possible explanation is that the game was only launched three months before our data were collected and the game developer was probably still learning how to profitably set prices (Dube and Misra, 2019; Huang et al., 2020). Second, while more flexible and discriminatory pricing strategies would lead to larger profit, the relative gains are limited when compared to simple uniform pricing: first-degree price discrimination would generate a mere 2% increase in profit over (optimal) uniform pricing. Our analysis suggests that this is the result of unsophisticated and myopic player behavior, which limits the extra gains of more elaborate pricing strategies. Third, each of the alternative pricing strategies considered—including uniform pricing—would induce a transfer of surplus from the players to the game developer without, however, generating sizeable dead-weight losses on average.

Our first finding is consistent with the aforementioned literature documenting that, sometimes, for-profit firms do not maximize profits (Cho and Rust, 2010; DellaVigna and Gentzkow, 2019; Dube and Misra, 2019; Fioretti, 2020; Hortaçsu et al., 2021; Huang et al., 2020; Orbach and Einav, 2007). Our second result is in line with Chu et al. (2011), who show that in the context of a theater company, simple pricing rules can sometimes generate almost as much profit as complex ones that would however be difficult to implement. Our second result is also close in spirit to Levitt et al. (2016), who document limited gains of second-degree price discrimination for a large online gaming firm, and more in general to the empirical literature on the trade-offs of price discrimination and personalized pricing in the era of big data (Rossi et al., 1996; Shiller and Waldfogel, 2011; Shiller, 2015; Waldfogel, 2015). Limited gains from price discrimination may partly explain why it is rarely observed in business practice, where—as already mentioned above—additional risks tied to consumer backlash also need to be considered (Council of Economic Advisors, 2015; DellaVigna and Gentzkow, 2019).

In contrast to our results, however, Dube and Misra (2019) document substantial returns of personalized pricing for a digital recruiting firm, highlighting the need for caution in drawing general conclusions: while we do not see any evidence for this in our analysis, in other digital contexts more complex pricing strategies may be more profitable. That said, both our results and Dube and Misra (2019) stress the large potential of "empirical" pricing rules. In the case of the game we study, by optimally choosing a uniform price on the basis of detailed data and appropriate empirical methods, the game developer could increase profit more than fourfold. Importantly, our third insight stresses that, although these increases in profit would necessarily come at the expense of consumer surplus, the pricing strategies considered would not generate average losses in total welfare.

Our paper contributes to a recent and growing literature investigating various aspects of mobile apps. Due to data limitations, most researchers have either exclusively focused on the supply side or employed very aggregate measures of demand, such as aggregate rankings or number of downloads from app stores (Bresnahan et al., 2015; Carare, 2012; Ershov, 2018; Ghose and Han, 2014b; Yi et al., 2019; Yin et al., 2014; Yuan, 2020; Wen and Zhu, 2019). Our user-level panel data instead allow us to delve deeper into the in-app purchase behavior of about 300,000 users around the world and to investigate the efficiency of discriminatory pricing strategies in a mobile game.

Despite widespread interest amongst practitioners and scholars alike (Fudenberg and Villas-Boas, 2006, 2012; Varian, 1989), there is relatively limited empirical evidence on the returns of price discrimination in practical applications, and essentially none for mobile games.[2] In

[2] Even though we focus on first-degree and third-degree price discrimination, there is a small empirical literature investigating the returns of second-degree price discrimination (quantity discounts): in carbonated soft drinks (Iaria and Wang, 2021), in coffee shops (McManus, 2007), in cable television (Crawford and Shum, 2007), in the yellow pages (Aryal and Gabrielli, 2020), and for an online gaming company (Levitt et al., 2016).

general, the extant empirical evidence is mixed, documenting limited returns in some cases (Rossi et al., 1996; Levitt et al., 2016; Shiller and Waldfogel, 2011; Shiller, 2015; Waldfogel, 2015) but larger in others (Adams and Williams, 2019; Cho and Rust, 2010; DellaVigna and Gentzkow, 2019; Dube and Misra, 2019; Iaria and Wang, 2021; List, 2004). Our paper contributes to this debate by providing the first empirical investigation on price discrimination for a mobile game. Despite the focus on a specific game, our empirical analysis speaks to a broader audience than freemium game providers: pay-gates are important monetization mechanisms also for other types of digital content providers including newspapers, magazines, and streaming services (e.g., Amazon Prime and You Tube).[3]

From a methodological perspective, ours is one of few empirical papers that combine both structural methods and randomized experiments (Cohen et al., 2016; Dube and Misra, 2019; Einav and Levin, 2010; Levitt and List, 2009; Todd and Wolpin, 2020). Our structural demand model for game content is needed to simulate the likely welfare consequences of counterfactual pricing strategies not observed in the data, while the experimental variation allows us to mitigate some of the standard endogeneity issues that would otherwise cripple identification and estimation.

The paper continues as follows. Section 2 describes the game for which we have data and discusses the way prices are set. Section 3 describes the data and the available sources of exogenous variation. Section 4 describes how we model player behavior and Section 5 reports our estimation results. Section 6 discusses our simulations for a number of counterfactual pricing rules. Section 7 draws some conclusions.

2 Mobile Game

2.1 Game Description

We empirically investigate the efficiency of price discrimination in the context of a game app which was produced by a large mobile game developer ("firm") and launched in August 2013. Like other popular mobile games, such as Candy Crush Saga or Bejeweled, the game we study is a casual game characterized by a sequence of levels that can be cleared in a relatively short amount of time. It belongs to the mobile puzzle game genre and has been downloaded around 80 million times so far, making it one of the most popular match-3 games of all time. The goal for players is to clear levels by connecting lines of jellies of the same color in order to "splash" them and achieve varying objectives.

[3]Another related emerging literature is that on algorithmic pricing. The majority of economics papers in this literature have so far been theoretical, mainly about the potential for algorithmic pricing to facilitate collusion (Miklós-Thal and Tucker, 2019; Calvano et al., 2020; Brown and MacKay, 2021), even though a few studies in progress are investigating the topic empirically, such as Assad et al. (2020) for gas stations and Hortaçsu et al. (2021) for airline companies.

Figure 1: 3-match mobile game

The initial allocation of jellies is random and a move consists of connecting at least three jellies of the same color; the longer the line (also called "snake") of connected jellies, the more points are awarded (see Figure 1). Connected jellies are removed and replaced by a random set of new jellies. Players must achieve different objectives to clear different levels, for example reach a minimum score, remove slime, move diamonds from top to bottom, and so forth, which are all achieved by connecting and removing jellies. The number of moves for each level is capped, and the maximum number of allowed moves varies by level. In contrast to traditional video games, the difficulty of each level does not increase as players advance. There are occasional spikes in difficulty in certain levels, although these do not occur at regular intervals (Debeauvais and Lopes, 2015). Levels distinguish themselves by their layout, objectives, or features, such as the presence of obstacles, and so on. To advance, players must clear every level. Once a level has been cleared, it can be replayed at any time.

Players are awarded a score for their performance, which largely depends on the length of the snakes formed as well as the total number of moves needed to clear a level. Upon clearing a level, players are awarded one, two, or three "stars" depending on their score for that level. Stars are cumulative and, as we explain below, play an important role in the monetization of the game, whereas the score is specific to each level and plays no role other than to determine the number of stars.

2.2 In-App Purchases and Monetization

The game is a freemium product. A certain number of levels can be played for free (with a few restrictions), but premium content, such as additional levels or features, need to be unlocked via in-app purchases. Importantly for our empirical analysis, during the period of our data collection, no in-app advertisement was displayed to players. This allows us to focus on in-app purchases as the only source of revenue for the firm and, in turn, to estimate a tractable choice model useful for the simulation of counterfactual pricing strategies. The co-existence of in-app purchases and advertisement would introduce dynamic interactions between pricing and advertising decisions which would be extremely hard to model, estimate, and ultimately simulate in counterfactual scenarios (Dubois et al., 2017).

In-app purchases must be paid in "virtual coins" and each player receives an initial endowment of 70 of these. This endowment of virtual coins corresponds to approximately $1 at the time our data collection. Once players have spent their endowment, they must purchase additional virtual coins to buy any of the following features.[4] First, players can purchase additional "moves" if they run out of these before having successfully cleared a given level. Second, players are initially endowed with five "lives." A life is lost every time a player attempts to but does not successfully clear a level. Lives replenish automatically, a life being added every 30 minutes. If a player loses all five lives, they either wait for 30 minutes before they can continue to play, or purchase a bundle of five lives. Alternatively, a player can gain lives by inviting friends to download the game via Facebook.

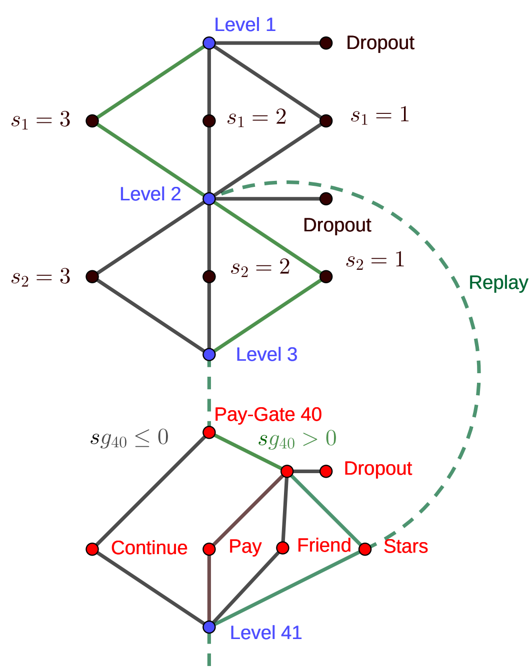

Figure 2: Players' decision tree

Figure 2 illustrates player i's decision tree. They start at level one and then either stop playing immediately or clear this level. By clearing any level, i can obtain three, two, or one star depending on the score obtained in that level. Starting from level 40, every 20 levels player i meets a "pay-gate," which we denote by $t = 40, 60, ..., T$ (i.e., pay-gates appear

[4] Note that, since our data were collected, some of the game's features have changed and additional opportunities for in-app purchases were introduced.

at level 40, 60, etc.). These pay-gates separate free levels from premium levels, and must be unlocked for i to proceed in the game. Players have three options to unlock pay-gates (see Figure 3): (a) purchase a "key" using 70 virtual coins, (b) invite friends on Facebook to download the game, or (c) accumulate a sufficient number of stars.

Figure 3: Example of a pay-gate.

Regarding option (c), each pay-gate has a threshold number of stars that is pay-gate specific and rising as the player progresses through the game. If a player gets to the pay-gate with a number of stars equal or greater than this threshold, the pay-gate unlocks. For brevity, we refer to the difference between i's number of accumulated stars through their play up to pay-gate t and the number of stars needed to unlock pay-gate t as "star gap" and we denote it by $sg_{i,t}$. Only players with a positive star gap ($sg_{i,t} > 0$) must unlock pay-gate t. To do so, they can use either option (a) or (b), or alternatively can go back and re-play previous levels to gain additional stars where they obtained less than three, a behavior called "grinding."

At the moment of the data collection, the firm was relying on a simple uniform pricing strategy of 70 virtual coins (approximately \$1) across all pay-gates and players for the purchase of a key to unlock a pay-gate.[5] As illustrated in Figure 4, the purchase of keys to unlock the first three pay-gates (levels $40, 60, 80$) corresponded to the largest share of in-app purchases (43%). Because of this and to maintain the econometric model and simulations practically viable, we focus on the firm's choice of which prices to charge for the keys to unlock these

[5]Some form of price discrimination was still implemented by offering features in bundles as well as by offering quantity discounts on larger amounts of virtual coins.

three pay-gates.[6]

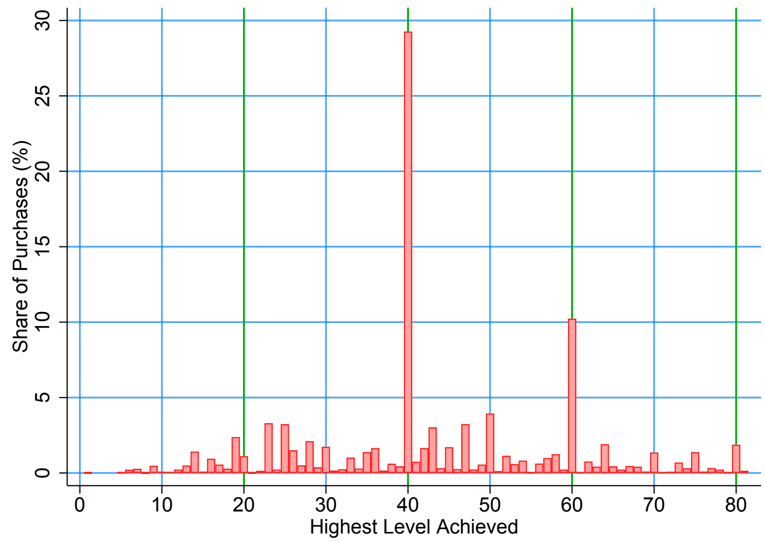

Notes: This figure displays the relative frequency of purchases observed in the game according to the highest level achieved. For example, a player who reaches pay-gate 40 with a positive star gap has as highest level achieved $\ell = 40$. If keys are purchased at this pay-gate, we count a purchase for $\ell = 40$. The sample includes all players (i.e, Group 20, Group 40, and Group No Star) and concerns only purchases at or before reaching level $\ell = 80$.

Figure 4: Spending patterns across life cycle

3 Data

3.1 Data and Variables

We have tracking data for all users around the world that installed the game between October 30$^{\text{th}}$ and November 4$^{\text{th}}$ 2013 on Apple devices (iPhones and iPads) and that played at least one round of the game. We have a sample of 292,179 players, and for each we observe the full history of play at an extraordinary level of detail for the 15 days following the installation of the game, including any purchase of virtual coins.[7] In particular, we rely on the following information to describe players' behavior and characteristics.

Level attempted and completed: For each player, we observe the level played in any given round of playing.[8] This allows us to track the sequence in which different levels are played and re-played. We also observe whether a level was cleared or not at a given attempt. Finally, we assume that a player drops out after the last attempt to clear a level.

[6]As discussed below, in Appendix A.2 we also provide supporting evidence that purchases of keys do not appear to crowd out other in-app purchases, suggesting that the two can be studied separately without excessive loss of generality.

[7]Parts of these data were also used in Wagner and Runge (2018).

[8]Rounds denominate the cumulative number of levels played irrespective of whether a level was played multiple times or not.

Score and stars awarded: We observe the score each player was awarded for clearing a level. As discussed above, the score reflects how well a player performs in a given level. The number of stars awarded in a level is then determined as a function of the score obtained (stars are only awarded if a level is successfully cleared and star thresholds vary across levels).

Player's ability: A unique feature of online games, as opposed to more traditional offline games, is the possibility of measuring, almost in real time, a player's gaming ability. We measure a player's ability as the average snake-length over the first 20 rounds played, where the longer the average snake of connected jellies, the larger the score obtained by the player. The larger our measure of ability, the greater the player's skill in identifying the patterns required to succeed in the game. As shown in Figure 5, this measure of player's ability is almost normally distributed with some additional concentration around the mean. It can also be interpreted in relative terms: the top player (with a measure of 11.6) can be considered "twice as able as" the average player (with average measure of 5.4).

Controlling for ability is important both economically and econometrically. Economically, willingness to pay may vary with players' ability, affecting the optimal pricing strategy of the firm. Analogously, not controlling for players' ability may give rise to complex forms of endogeneity hard to address econometrically.

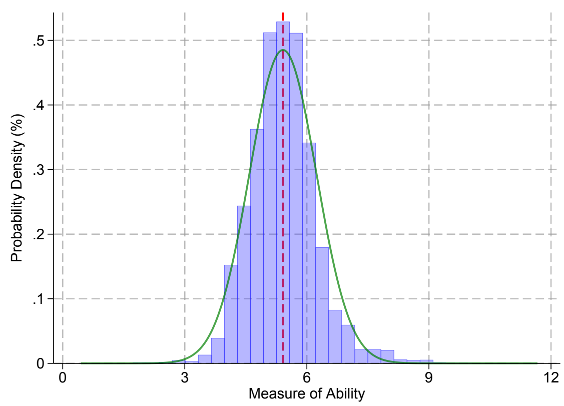

Notes: This figure displays the distribution of our measure of ability across players. The mean is presented through the vertical red line and the plot is overlaid with a normal distribution in green. The description of this variable is provided in Section 3.1. The sample is a cross-section of all players (Group 20, Group 40, and Group No Grind).

Figure 5: Distribution of players' ability

Star gap: As discussed above, each pay-gate is unlocked (and will not appear again) if a player reaches it with a sufficient number of stars, what we call a non-positive star gap. When a player reaches a locked pay-gate (i.e., with a positive star gap), they cannot proceed in the game until they unlock it. When reaching a locked pay-gate, players can grind in an attempt to decrease their star gap enough so to unlock the pay-gate. To simplify the analysis, we consider a measure of star gap inclusive (or gross) of grinding, rather than considering grinding as a separate decision: we measure $sg_{i,t}$ to be i's star gap at pay-gate t *after* all the grinding—when they either unlock t or drop out of the game. If star gaps were measured *net* of grinding, i's decision at pay-gate t with $sg_{i,t} > 0$ would be whether to unlock t, to grind in order to lower $sg_{i,t}$, or to stop playing. Differently, by measuring star gaps *gross* of grinding, we simplify i's decision at pay-gate t with $sg_{i,t} > 0$ to be only between unlocking t or stopping to play, given that $sg_{i,t} > 0$ is already inclusive of all of i's grinding at t.[9]

Pay-gate locked and unlock mechanism: We observe whether a player reaches a locked pay-gate and how they unlock it (or whether they drop out of the game). As mentioned above, when approaching pay-gate t with fewer stars than those necessary to unlock it, $sg_{i,t} > 0$, player i's options to unlock t are: (a) paying 70 virtual coins to purchase a key, (b) inviting a friend on Facebook to download the game, or (c) going back to previous levels to collect more stars (i.e., grinding). Importantly for our econometric analysis, as discussed in detail below, players are not notified about the appearance of pay-gates every 20 levels when they start playing.

Price to unlock a gate: Any player i reaching pay-gate t with a positive star gap $sg_{i,t} > 0$ cannot proceed in the game without unlocking it. As discussed above, a way to unlock a pay-gate is to purchase a key. In the period of our data, the price of a key was set by the firm to $p_{i,t} = 70$ virtual coins, approximately $1, uniformly for any i and t. The appropriate choice of $p_{i,t}$ (in terms of virtual coins) by the firm, potentially discriminating across i's and t's, is the main object of our empirical analysis. Because keys are priced in virtual coins and players can rely on their endowment of virtual coins to buy keys (thus spending potentially less in terms of real money), we consider $p_{i,t}$ as the effective or residual price of purchasing a key: e.g., the full price of the key minus i's endowment of virtual coins when reaching t (e.g., if the full price of a key is 70 virtual coins and i owns 30 virtual coins when reaching pay-gate t with $sg_{i,t} > 0$, then $p_{i,t} = 70 - 30 = 40$).

Player demographics: We observe a number of player-specific characteristics measured when a player downloads the game. These variables are collected in a vector we call X_i throughout the paper. Most of these characteristics relate to the device used to play the game. We know whether the game was downloaded to a mobile phone or a tablet (iPad). We also observe whether a player has updated their device to the latest version of the relevant operating system (iOS7) and whether the device was "jailbroken" by its owner (Jailbroken).[10] Finally, we observe the country of a player (as indicated by the national app store used to download the app) and relate it to its 2013 GDP per capita measured in purchasing power parity. We assign these countries to fourteen groups which we refer to as "regions." This assignment is detailed in Appendix G and the share of players in each region is displayed in Figure A-26.

[9]This will become clearer after having formally specified the choice model in Section 5, see in particular footnote 15 and the surrounding discussion.

[10]Jailbreaking means removing all restrictions imposed on the device in order to allow the installation of software not supported by Apple.

Table 1 reports summary statistics for X_i among the 292,179 players in the data.

	mean	sd	min	max
Maximum Level Reached	21.30	17.16	0.00	179.00
Player's Ability	5.41	0.82	0.44	11.67
Log(GDP per Capita (PPP, 2013))	10.55	0.48	6.54	11.85
Jailbroken Dummy	0.01	0.11	0.00	1.00
iOS7 Dummy	0.78	0.42	0.00	1.00
iPad Dummy	0.31	0.46	0.00	1.00
Num. of Players	292,179			

Notes: This table provides descriptive statistics for the demographic variables. The definitions of the variables are detailed in Section 3.1. The sample includes all players (Group 20, Group 40, and Group No Grind) and the statistics are computed across this cross-section of players.

Table 1: Descriptive statistics of players' characteristics X_i

3.2 Exogenous Variation

In addition to the extremely detailed player-specific information described above, our data are also unique in providing various sources of exogeneous variation helpful to characterize players' behavior and to identify our econometric model. The first source of exogeneous variation is represented by controlled experiments conducted by the firm during the period of our data collection. The second is represented by a form of randomness in the degree of difficulty faced by different players when playing any level.

3.2.1 Controlled Experiments

During the period of our data collection, the firm conducted a controlled experiment that randomly allocated players to three different designs of the pay-gates separating free from premium levels. Figure 3 illustrates what the firm considered the default design of the pay-gates in the game: 40 free levels before the first pay-gate appears, with a new pay-gate appearing every 20 levels thereafter. The default design allows for three options to unlock a pay-gate: (a) paying 70 virtual coins to purchase a key, (b) inviting friends to download the game via Facebook, or (c) having a non-positive star gap $sg_{i,t} \leq 0$. About 16% of all players were allocated to this default design (called *Group 40*).

The experimental variation introduced by the firm consists of two variations relative to the default design. In the first variation (called *Group 20*), a subset of nearly 16% of players was exposed to an earlier first pay-gate already after clearing level 20. This setting allows for the same three options to unlock pay-gates as the default design. In the second variation (called *No Stars*), the first pay-gate appears after clearing level 40 as in the default design, however option (c) to unlock pay-gates with non-positive star gaps is not available. When i from this group reaches any pay-gate t, independently of $sg_{i,t}$, they must choose either option (a) or

(b) to unlock it and proceed in the game. The No Stars group represents around 68% of players in our sample. Table 2 summarizes the main features of these three groups.

Group	**Level 1st pay-gate**	**Unlock with sg≤0**	**Share in sample**
Group 40	40	Yes	0.16
Group 20	20	Yes	0.16
No Stars	40	No	0.68

Table 2: Experimental groups

Table A-15 in the Appendix provides descriptive statistics for these three groups, confirming that player-specific characteristics are well-balanced.

Player attrition. Player attrition is a common feature in this type of game. Figure 3.2.1 shows the share of active players at the start of each level (i.e., players who have not yet dropped out), comparing Group 20 to Group 40 (left panel) and No Stars to Group 40 (right panel). The graph shows that attrition is high in all three groups as players progress in the game. Overall, only about 24% of all players clear level 40 and reach pay-gate 40. Not surprisingly, the left panel of Figure 3.2.1 reveals a slight difference in attrition rates at level 20 between Group 20 and Group 40, as only players in Group 20 face a pay-gate at level 20. However, between levels 21 and 40, the gap between the two groups closes again and approximately 23.9% of Group 20 reaches level 40 compared to 24.5% of Group 40. The right panel of Figure 3.2.1 shows similar attrition rates for Group 40 and the No Stars group at level 40, suggesting that both groups are strongly affected by the appearance of the first pay-gate. Overall, this evidence is in line with Debeauvais and Lopes (2015), who document—for a different cohort of players—that attrition is larger in levels with pay-gates than in regular levels. In our analysis, we discard observations beyond level 80. This is essentially without loss of generality, in that 99.95% of all players drop out before level 80.

Choices at pay-gates. In Table 3, we summarize players' choices at the pay-gates of levels 20 and 40. In Group 20 about 83.09% of all players reached the first pay-gate at level 20 with sufficient stars to unlock it ($sg \leq 0$). The vast majority of players without sufficient stars ($sg > 0$) used the initial endowment of virtual coins to purchase a key (13.93%), only 1.41% used real money, while 1.56% invited friends on Facebook to download the game.

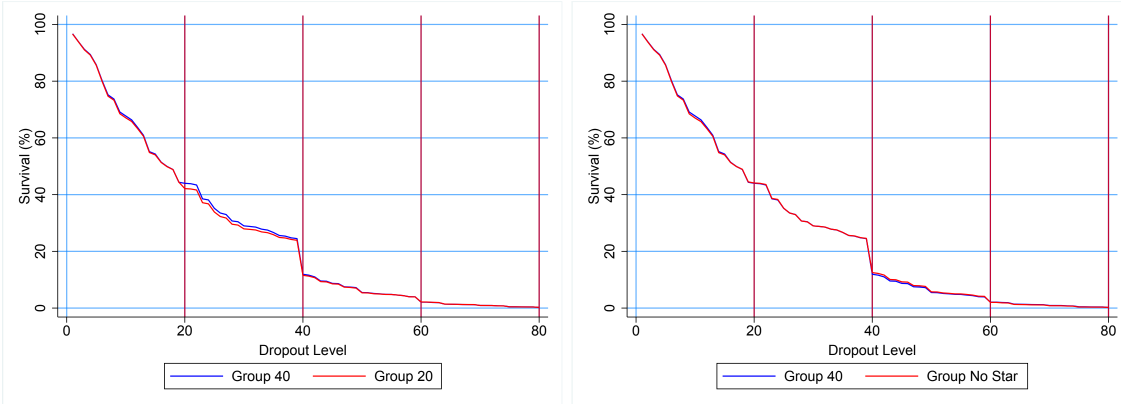

Notes: This figure displays the share of players according to the highest level ℓ they have reached, by treatment group. The left panel compares Group 20 and 40. The right panel compares Group 40 and Group No Star. These treatment groups are defined in Section 3.2.1. The sample includes all players.

Figure 6: Kaplan-Meier survival function (by group)

		Group 20		**Group 40**		**No Stars**	
		1st pay-gate at level 20		1st pay-gate at level 40			
		Number	%	Number	%	Number	%
Players reach pay-gate 20		19,183	44.38	19,367	44.35	91,394	44.49
Players unlock pay-gate 20		18,203	94.89				
Unlock option							
$sg \leq 0$		15,126	83.09				
$sg > 0$	Buy key: Real money	257	1.41				
	Buy key: Endowment	2,536	13.93				
	Facebook	284	1.56				
Players reach pay-gate 40		10,329	23.90	10,694	24.49	50,469	24.57
Players unlock pay-gate 40		4,989	48.30	5,203	48.65	25,701	50.92
Unlock option							
$sg \leq 0$		1,135	22.75	1,094	21.02		
$sg > 0$	Buy key: Real money	853	17.10	826	15.87	5,350	20.81
	Buy key: Endowment	1,791	35.89	2,168	41.66	14,121	54.94
	Facebook	1,210	24.25	1,115	21.42	6,230	24.24

Table 3: Comparison of unlock mechanisms between treatment and control groups

This indicates that the pay-gate at level 20 is a relatively soft monetization trigger, as most players were able to unlock it either with $sg \leq 0$ or using their initial endowment of virtual coins. At level 40, the share of players with $sg \leq 0$ is significantly lower compared to the pay-gate at level 20, yet similar across Group 20 (22.75%) and Group 40 (21.05%). The share of players purchasing a key using real money increases to 17.10% for Group 20 and to 15.87% for Group 40. As expected, players in the No Stars group, which cannot unlock pay-gate 40 using their accumulated stars, were significantly more likely than the others to purchase a key, either using real money (20.81%) or their endowment of virtual coins (54%).

As discussed in Sections 4 and 5, we use this experimental variation to test competing hypotheses of players' behaviour and specify a more appropriate choice model, such as the degree of forward-looking behavior with respect to upcoming pay-gates (comparing Group

20 and Group 40), and to overcome identification concerns about endogenous selection on positive star gaps in estimation (relying on the No Stars group).

3.2.2 Randomness in the Difficulty of Levels

The structure of the game offers another useful source of exogenous variation: the difficulty of levels. Every time a new round of the game is played, there is a random draw of jellies which may incidentally deliver an easier or harder problem for the player to solve. A "good" draw may lead the player to succeed at a given level, while a "bad" one may be enough to induce the same player to fail. When a player gets closer to failing a level, they may face more of an incentive to purchase and spend virtual coins to obtain additional lives or moves, so to get the final boost needed to clear the level. Controlling for a player's ability, worse random draws of jellies will result in stronger incentives to purchase and spend virtual coins for reasons other than a key.

Following this line of reasoning, as detailed in Section 5, we exploit i's "bad luck in the random draws of jellies" as a source of exogenous variation—an instrument—for i's effective price of a key $p_{i,t}$ (defined as 70 virtual coins minus i's endowment) in the estimation of i's probability to purchase a key.

4 Characterizing Player Behavior Using Exogenous Variation

4.1 The Firm's Revenue Function

Our goal is to estimate a realistic but parsimonious model of player behavior useful to simulate alternative pricing strategies and their returns both to the firm and to players. We focus on the firm's revenue from purchases of keys and ignore other in-app purchases (e.g., additional lives or moves). This approach is motivated by the evidence shown in Figure 4 (most in-app purchases are concentrated at pay-gates) and the absence of any in-app advertising at the time of our data collection. We also show in Appendix A.2 that purchases of keys do not crowd out other in-app purchases, suggesting that the two can be studied separately without excessive loss of generality.

We aggregate player i's decisions between any two pay-gates into a single choice, abstracting from the intermediate choices made at each level. Our choice model has two components. First, the probability that i reaches pay-gate t (denoted by $i \to t$) with a positive star gap given that she already unlocked pay-gate $t - 20$:

$$Pr_{i,t}(i \to t, sg_{i,t} > 0 | t + 20, ..., T). \tag{1}$$

Second, conditional on pay-gate t being locked to i (denoted by the indicator $\text{lock}_{i,t} = 1$) and effective price $p_{i,t}$, the probability of purchasing a key at t (denoted by $\text{buy}_{i,t} = 1$, we describe this categorical variable in Section 5 in more detail):

$$Pr_{i,t}(\text{buy}_{i,t} = 1 | \text{lock}_{i,t} = 1, p_{i,t}, t + 20, ..., T). \tag{2}$$

Relying on (1) and (2), we then specify the firm's expected revenue from player i at pay-gate t, $R_{i,t}$. Finally, our simulation exercises entail the maximization of $R_{i,t}$ with respect to $p_{i,t}$ across all players and pay-gates—under various constraints on the flexibility of prices (from uniform price to first degree price discrimination). Because every additional player does not generate any increase in the firm's costs (at least within the range observed in our sample), throughout the paper we assume that the firm's marginal costs are zero, and that expected revenue equals expected profit.

In general, probabilities (1) and (2) could be complex functions of i's expectations about future realizations of any variable (e.g., $sg_{i,t+20}$ and $p_{i,t+40}$). To keep the empirical model manageable, especially in view of our extensive simulation exercises, we propose (and then verify empirically) the following simplifying assumption.

Assumption 1 (Myopia)**.** *Players' decisions in t are conditionally independent of expectations about future decisions and variables to be realized in $t + 20, t + 40, ..., T$.*

Assumption 1 implies that player behavior can be represented by a (i, t)-specific static choice model, so that (1) and (2) simplify to $Pr_{i,t}(i \to t, sg_{i,t} > 0)$ and $Pr_{i,t}(\text{buy}_{i,t} = 1 | \text{lock}_{i,t} = 1, p_{i,t})$, respectively. Importantly, the conditional independence implied by this assumption should be intended with respected to the observable characteristics we can control for in the data, such as i's demographics and ability, and various fixed effects. We discuss the details of the empirical specification of our model in Section 5.

Denote i's effective prices from pay-gate $t + 20$ until T by $p_{i,>t} = (p_{i,t+20}, p_{i,t+40}, ..., p_{i,T})$. Then, given choice models (1) and (2) and Assumption 1, the firm's expected revenue from player i at pay-gate t of charging effective prices $p_{i,\geq t}$, given that i already unlocked pay-gate $t - 20$, can be expressed as:

$$\begin{aligned}
&R_{i,t}(p_{i,t} | p_{i,>t}) \\
&= Pr_{i,t}(i \to t, sg_{i,t} > 0) \times Pr_{i,t}(\text{buy}_{i,t} = 1 | \text{lock}_{i,t} = 1, p_{i,t}) \times p_{i,t} \\
&+ Pr_{i,t}(i \to t, sg_{i,t} > 0) \times \big(1 - Pr_{i,t}(\text{buy}_{i,t} = 0 | \text{lock}_{i,t} = 1, p_{i,t})\big) \times R_{i,t+20}(p_{i,t+20} | p_{i,>t+20}) \\
&+ Pr_{i,t}(i \to t, sg_{i,t} \leq 0) \times R_{i,t+20}(p_{i,t+20} | p_{i,>t+20}).
\end{aligned} \tag{3}$$

As this expression indicates, despite the simplifying assumptions, the firm's expected revenue from i at pay-gate t is an intricate recursive function. Conditional on i having unlocked pay-gate $t-20$, it depends on i's probability of reaching pay-gate t with $sg_{i,t} > 0$ (so that pay-gate t is locked, $\text{lock}_{i,t} = 1$), i's probability of purchasing a key ($\text{buy}_{i,t} = 1$) given $\text{lock}_{i,t} = 1$, i's probability of unlocking the current pay-gate ($\text{buy}_{i,t} \neq 0$) given $\text{lock}_{i,t} = 1$, price $p_{i,t}$, and finally—should i unlock the current pay-gate—the expected revenue stemming from potential purchases of keys to unlock future pay-gates.[11] The first line of model (3) denotes the firm's expected revenue from i's *current* purchase of a key to unlock pay-gate t, the second and third lines instead denote the firm's expected revenue from i's *future* purchases of keys to unlock pay-gates $t + 20, t + 40, ..., T$. Before getting to the details of how we specify the components of model (3), we rely on the exogeneous variation described above to test the consistency of Assumption 1 with observed player behavior.

4.2 Testing Assumption 1

The validity of the expected revenue function in (3) crucially depends on the validity of Assumption 1 for the game we study. Here we rely on the exogenous variation available in the data to provide empirical evidence in support of this simplifying assumption.

Assumptions 1 requires that players' current choices are not influenced by their expectations regarding future events. The experimental variation in our data allows us to test for the absence of forward-looking behavior in various ways. In particular, we exploit the exogenous information shock to players in Group 20, who become aware of the existence of pay-gates twenty levels before the other players.

If players were forward-looking, those in Group 20 could show different attrition rates compared to other players. Having passed the the first pay-gate at level 20, their expected utility from continuing the game could be lower due to the anticipation of additional pay-gates (i.e., entailing costs with positive probability) at future levels. As a consequence, attrition rates could be higher. In particular, we test for differences in the total number of rounds played between players in Group 20 that were exposed to the pay-gate at level 20 but had $sg_{i,t} \leq 0$ and similar players in Group 40 who were not exposed to the pay-gate at level 20. These "similar" players in Group 40 are those who had a sufficient number of stars to immediately unlock the pay-gate at level 20 had they been allocated to treatment Group 20. We do not find significant differences (Table 4, row 1).

[11] As mentioned in Section 2.2, players can unlock pay-gates not only by purchasing keys but also by asking friends on Facebook to download the game. In this sense, as detailed in Section 5, the categorical variable $\text{buy}_{i,t}$ can take more values than only 0 (stop playing) and 1 (purchase a key), and $(1 - Pr_{i,t}(\text{buy}_{i,t} = 0|\text{lock}_{i,t} = 1, p_{i,t})) \geq Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t})$.

	Group 40	Group 20	Diff.	Std. Err.	Obs.
Number of rounds before drop out	121.026	122.860	-1.833	1.633	21,899
Rounds played between 21 and 40 (or drop out)	51.400	51.582	-0.181	0.509	21,899
Stars collected between 21 and 40 (or drop out)	24.014	23.951	0.063	0.149	21,899
Re-played levels between 21 and 40 (or drop out)	26.375	27.354	-0.978	0.761	21,899

Notes: This table presents evidence regarding players' forward-looking behavior. The description of the variables is provided in Section 3.1. The sample of players includes all players in Group 20 and 40 who have crossed level $\ell = 20$ with a non-positive star gap. Columns "Group 40" and "Group 20" report the mean for Group 40 and 20 players, respectively. Column "Diff." provides the (mean) difference between the two former columns. Column "Std. Err." presents the standard errors associated with the mean of column "Diff.".

Table 4: Experimental evidence for myopia: Group 20 vs Group 40

Awareness of the existence of pay-gates should affect the propensity to grind of forward-looking players (i.e., re-play past levels to collect additional stars). A player that is aware of the existence of future pay-gates, and the possibility of unlocking them with a sufficient number of stars, should grind more than unaware players, in order to increase the chance of reaching the next pay-gate with a non-positive star gap. However, we find no significant difference in the number of rounds played between levels 21 and 40 for the same groups of players used in the previous test (Table 4, row 2). We also find no significant differences in the number of stars collected (Table 4, row 3) and in the number of re-played levels between levels 21 and 40 (Table 4, row 4).[12]

We look for evidence of forward-looking behavior in two additional ways. First, we test for evidence of forward-looking behavior by checking whether the number of additional stars collected after having cleared a level for the first time affects the player's probability to re-play that level. For example, if player i cleared level 23 for the first time with $n \in \{1, 2, 3\}$ stars and is able to move on to level 24, we check if their probability to re-play level 23 depends on n. Obtaining only one or two stars leaves open the possibility to collect an additional two or one stars, respectively, by re-playing the level, hence increasing the chance of reaching the next pay-gate with a non-positive star gap. Table 5 shows estimation results for a multinomial logit model of the probability to re-play any level between 21 and 40 that was cleared for the first time with n stars (with the case of $n = 3$ stars as the excluded category, i.e. no incentives to re-play those levels). In line with economic intuition, the estimated intercepts suggest that the probability to re-play any level in 21-40 is increasing in the number of stars a player can still collect by re-playing it. The Group 20 indicator, however, is not significantly different from 0, providing no statistical evidence in support of forward-looking behavior. This is also shown graphically in a more disaggregate way, level by level starting from level 1, in Appendix Figures A-1, A-2, and A-3. These graphs confirm that players in Groups 20 and 40 have virtually identical probabilities of re-playing any specific level initially cleared with a given number of stars, both before and—importantly—after players in Group 20 become aware of the existence of the pay-gate at level 20.

[12]In additional tests (not reported, but available on request), we also check the propensity to re-play specific levels for the same group of players and find no significant differences.

	Multinomial Logit
$Pr(\text{Re-play level in 21-40}, \text{Stars} = 1)$	
Group 20 (relative to Group 40)	0.014
	(0.057)
Constant	0.602
	(0.039)
$Pr(\text{Re-play level in 21-40}, \text{Stars} = 2)$	
Group 20 (relative to Group 40)	0.013
	(0.053)
Constant	0.083
	(0.037)
Observations	20,997

Notes: This table presents the results from multinomial logit regressions where we estimate the propensity to replay levels depending on the number of stars collected when first clearing the level. Each observation is a case where a player re-played a level of the game which was previously cleared. The sample includes all players in Group 40 and 20 which cleared level $\ell = 20$ with a non-positive star gap. The base category is $(\text{Re-play level in 21-40}, \text{Stars} = 3)$. Standard Errors are clustered at the player-level.

Table 5: Experimental evidence for myopia: Group 20 vs Group 40

Second, we inspect the distribution of effective prices players face when they reach the pay-gate at level 40. Remember that $p_{i,t}$ is defined as 70 virtual coins (the price of a key to unlock any pay-gate) minus i's residual endowment of virtual coins by the time they reach t. Since the initial endowment provided to every player is 70 virtual coins, by not using any of it until level 40, i would face $p_{i,40} = 0$ and be able to purchase a key to unlock pay-gate $t = 40$ without spending any real money. However, players can spend their endowment on a number of items before reaching pay-gate $t = 40$ to enhance their play experience, such as purchasing boosters, additional lives or moves, etc. Awareness of the existence of pay-gates is expected to induce forward-looking players to save up on their endowment of virtual coins, so to ensure a lower effective price at the next pay-gate. We then ask whether players in Group 20 with $sg_{i,20} \leq 0$, who are aware of the existence of pay-gates from level 20, spend their endowment between levels 21 and 40 differently than players in Group 40, who are unaware of the existence of pay-gates until level 40. Figure 7 illustrates the distribution of effective prices faced by players in Group 20 and Group 40 when reaching the pay-gate at level 40. The distributions of effective prices faced by the two groups is very similar and indicates that awareness of the existence of pay-gates did not lead players in Group 20 to save more of their endowment of virtual coins in anticipation of the next pay-gate. To corroborate these results, Appendix Figures A-4 and A-5 repeat the same exercise for pay-gates $t = 60$ and $t = 80$.

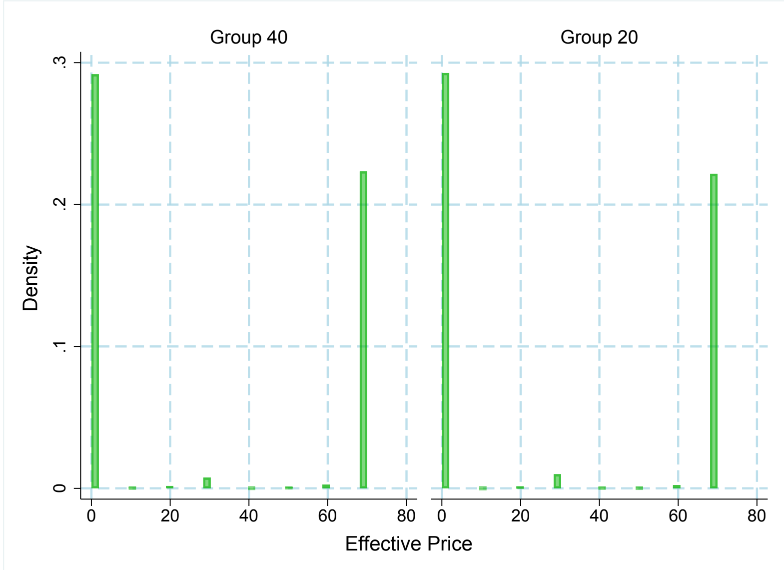

Notes: This histogram displays the effective price distribution at pay-gate 40 for Group 20 and Group 40. Each observation is a player who faces pay-gate 40 with a positive star gap. The sample includes all players of Group 40 and 20 who satisfy this condition.

Figure 7: Distribution of effective prices $p_{i,t}$ at pay-gate $t = 40$

5 Choice Model: Specification and Estimation

Our model abstracts from players' disaggregate level-specific choices and focuses on whether they reach each pay-gate t (discrete choice model (1)) and, conditional on reaching it with $sg_{i,t} > 0$, whether they choose to unlock it and how (discrete choice model (2)). Here we specify the empirical counterparts of these discrete choice models and estimate them relying on the exogenous variation described in Section 3.2.

5.1 Discrete Choice Model (1): Reaching Pay-Gates with a Positive Star Gap

We specify discrete choice model (1) as the product of two binary choice models:

$$Pr_{i,t}(i \rightarrow t, sg_{i,t} > 0) = Pr_t(i \rightarrow t|X_i) \times Pr_t(sg_{i,t} > 0|i \rightarrow t, X_i), \tag{4}$$

where X_i is a vector of observable i-specific characteristics such as i's demographics and player ability (see Section 3.1 for a description of these variables). After having unlocked pay-gate $t-20$, player i can either clear all levels between $t-19$ and t and reach the next pay-gate t, or stop playing before reaching it. This is the first binary choice model in (4), $Pr_t(i \rightarrow t|X_i)$. Upon reaching pay-gate t, we then distinguish between $sg_{i,t} > 0$ and $sg_{i,t} \leq 0$ to determine if i faces the next choice, discrete choice model (2), of whether and how to unlock pay-gate t. This is the second binary choice model in (4), $Pr_t(sg_{i,t} > 0|i \rightarrow t, X_i)$.

Neither of the binary choice models in our empirical specification (4) depends on the effective price $p_{i,t}$ of purchasing a key to unlock pay-gate t. While this is true by construction for $Pr_t(i \to t|X_i)$, in that $p_{i,t}$ can only be determined when i reaches pay-gate t, it may not be true for $Pr_t(sg_{i,t} > 0|i \to t, X_i)$. Given our definition of star gap as inclusive of grinding (see Section 3.1), this exclusion restriction may be violated for example if players were more likely to grind (and so to lower their star gaps) when facing higher effective prices, since unlocking pay-gates by non-positive star gaps would become relatively cheaper than by purchasing keys. Table 6 suggests this is not the case and reports supportive empirical evidence in favor of this exclusion restriction using the sample of players in Group 40,[13] a linear probability model for $Pr_t(sg_{i,t} > 0|i \to t, X_i)$ has an estimated coefficient on $p_{i,t}$ which is very close to zero, especially once we control for the observable i-specific characteristics X_i.

	(1) $Pr_t(sg_{i,t} > 0\|i \to t)$	(2) $Pr_t(sg_{i,t} > 0\|i \to t)$	(3) $\Pr_t(sg_{i,t} > 0\|i \to t, X_i)$
Effective price, $p_{i,t}$	0.0000359	0.0000515*	0.0000258
	(0.0000192)	(0.0000206)	(0.0000188)
Pay-gate fixed effects	No	Yes	Yes
Player-specific characteristics (X_i)	No	No	Yes
Observations	12,600	12,600	12,600
Num. of players	10,692	10,692	10,692

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table shows estimation results for the linear probability model of the effective price $p_{i,t}$ on player i facing a positive star gap at pay-gate t. The dependent variable is a dummy variable equal to one when a player has a positive star gap, as defined by equation (1), and zero otherwise. The explanatory variable is the effective price, defined in Section 3. The sample includes all players from Group 40 who have reached pay-gate 40, 60, or 80. Each observation is a player/pay-gate combination. In the first column, we include no controls. In the second column, we add pay-gate fixed effects. In the third column, we include the demographics X_i defined in Section 3.1. Standard errors are clustered at the player level.

Table 6: Linear Probability Model of Effective Price on Positive Star Gap

The exclusion of $p_{i,t}$ from the empirical specification of model (4) has practical implications for our analysis. Because of their conditional independence of $p_{i,t}$, the binary choice models in (4) will only act as "constant weights" in any of the maximizations of the firm's expected revenue (3) to be performed in our simulations. We therefore do not make any further assumption and estimate them as non-parametric functions of X_i for each pay-gate t.

We separately estimate each of the two binary choice models in (4) by a standard K-Nearest Neighbors (kNN) estimator from the sample of players in Group 40 (the same sample used in Table 6). We restrict the estimation of model (4) to the players in Group 40, since this is the default design of the game.[14] In Appendix B, we describe the kNN estimator of (4) and present its estimates, while in Figure 8 we plot the two estimated binary choice models as functions of player's ability.

[13] As explained below, we focus on the players in this group to estimate model (4).

[14] Players in Group 20 face an additional pay-gate at level 20 which increases their attrition, while those in No Stars cannot use their accumulated stars to unlock pay-gates; hence they have fewer incentives to obtain non-positive star gaps.

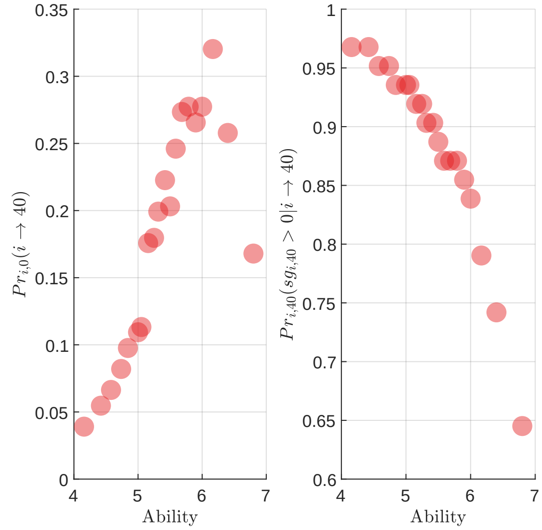

Notes: These figures display binned scatter plots of the relationship between the estimated probabilities of model (4) and ability. The estimates are produced using the kNN procedure described in Appendix B. Ability is defined in Section 3.1. The left panel displays the probability of reaching pay-gate 40 given that the player is at level 0, $Pr_{i,0}(i \to 40)$. The right panel displays the probability of a positive star gap given that the player has reached pay-gate 40 $Pr_{i,40}(sg_{i,40} > 0|i \to j)$. The sample used includes all players of Group 40.

Figure 8: kNN Estimates of Model (4) and Player's Ability

We see that, in line with intuition, more able players are less likely to drop out of the game (except for the top 5% of players) (left panel) and are more likely to reach pay-gates with a non-positive star gap (thus unlocking them without the need to purchase keys) (right panel).

5.2 Discrete Choice Model (2): Purchasing a Key to Unlock a Pay-Gate

The estimation of discrete choice model (2), $Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t})$, presents at least two challenges: sample selection on $sg_{i,t} > 0$ and endogenity of $p_{i,t}$.

First, given the structure of the game, only players with $sg_{i,t} > 0$ face $\text{lock}_{i,t} = 1$ and can be observed to purchase a key at pay-gate t. In the absence of experimental variation, we would then have to estimate (2) exclusively on the sample of players observed to reach pay-gate t with $sg_{i,t} > 0$, $Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t}, sg_{i,t} > 0)$. It is however possible that the willingness to purchase a key at t differs systematically between the players observed with $sg_{i,t} > 0$ and those observed with $sg_{i,t} \leq 0$, so that $Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t}, sg_{i,t} > 0) \neq Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t})$. Fortunately, the experimental variation described in Section 3.2 allows us to overcome this by restricting estimation of model (2) to the sample of players in the No Stars group. Players in the No Stars group cannot use their accumulated stars to unlock pay-gates and, independently of their observed star gap, always face $\text{lock}_{i,t} = 1$.

Second, the effective price $p_{i,t}$, computed as 70 virtual coins minus i's residual endowment of virtual coins at pay-gate t, depends on i's decision of whether to purchase and spend virtual coins before reaching pay-gate t. This decision, in turn, may correlate to i- and t-specific unobservable characteristics that also drive i's willingness to purchase a key at pay-gate t. We address this potential endogeneity in two ways. On the one hand, our empirical specification of model (2) controls for i's ability in the game, which would otherwise be the most worrying omitted variable. On the other, as mentioned in Section 3.2 and discussed in more detail below, we also exploit randomness in the difficulty of each level across players as an instrument for $p_{i,t}$: at the beginning of each level, different players get a random draw of jellies which determines the level's difficulty, and in turn the players' incentives to purchase and spend virtual coins for reasons other than a key (e.g., additional lives or moves).

Model Specification. When i reaches locked pay-gate t, $\text{lock}_{i,t} = 1$, for given effective price $p_{i,t}$, i faces three options, denoted by $\text{buy}_{i,t} \in \{0, 1, 2\}$:[15]

> $\text{buy}_{i,t} = 0$: Do not unlock pay-gate t and stop playing.
>
> $\text{buy}_{i,t} = 1$: Purchase a key to unlock pay-gate t at a price of $p_{i,t}$ virtual coins, where 70 virtual coins cost around \$1 in terms of real money.
>
> $\text{buy}_{i,t} = 2$: Ask a friend on Facebook to download the game.

In order for our simulations to be practically viable but still capture rich forms of observed and unobserved player-specific heterogeneity, we specify discrete choice model (2) parametrically as a mixed logit (McFadden and Train, 2000). Player i's conditional indirect utility from choosing $\text{buy}_{i,t}$ when facing locked pay-gate t is:

$$U_{\text{buy},i,t}(\eta_i) = \begin{cases} \epsilon_{0,i,t} & \text{if } \text{buy}_{i,t} = 0 \\ \delta_1 + \delta_{1,t} + \delta_{1,i} + X_i\beta_1 - (\alpha + \alpha_t + \alpha_i + X_i\pi)p_{i,t} + \epsilon_{1,i,t} & \text{if } \text{buy}_{i,t} = 1 \\ \delta_2 + \delta_{2,t} + \delta_{2,i} + X_i\beta_2 + \epsilon_{2,i,t} & \text{if } \text{buy}_{i,t} = 2, \end{cases} \tag{5}$$

where $(\delta_{\text{buy}} + \delta_{\text{buy},t} + \delta_{\text{buy},i})$ is an intercept given by the sum of a common δ_{buy} component among players, a pay-gate specific shift $\delta_{\text{buy},t}$ equal to zero at $t = 40$, and an unobserved player-specific random coefficient $\delta_{\text{buy},i}$, X_i is a vector of observable i-specific characteristics (see Section 3.1), $(\alpha+\alpha_t+\alpha_i+X_i\pi)$ denotes i's sensitivity to the effective price at pay-gate t, which is both a function of unobserved random coefficient α_i and observed heterogeneity $(\alpha+\alpha_t+X_i\pi)$, while $\epsilon_{\text{buy},i,t}$ is a residual error term we describe below. We allow for a particularly flexible specification of price sensitivity, itself a function of the observable characteristics X_i,

[15] Because our measure of star gaps is inclusive of grinding (see Section 3.1), any i with $sg_{i,t} > 0$ (and consequently with $\text{lock}_{i,t} = 1$) cannot—by definition—be observed to unlock pay-gate t by further grinding: all of i's grinding for pay-gate t is already included in $sg_{i,t}$. We therefore do not consider the option to grind as a further alternative to unlock pay-gates in model (10): all the grinding is captured by the second component of model (4), $Pr_t(sg_{i,t} > 0 | i \to t, X_i)$.

to investigate the potential of pricing strategies that take advantage of the detailed player-specific information routinely collected by the firm. We gather the three random coefficients $(\delta_{1,i}, \delta_{2,i}, \alpha_i)$ into the random vector η_i and assume that they are jointly normal:

$$\eta_i = \begin{pmatrix} \delta_{1,i} \\ \delta_{2,i} \\ \alpha_i \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho_{12} & \rho_{1\alpha} \\ - & \sigma_2^2 & \rho_{2\alpha} \\ - & - & \sigma_\alpha^2 \end{pmatrix}\right]. \tag{6}$$

As is standard, we normalize the systematic component of the indirect utility of $\text{buy}_{i,t} = 0$ (the outside option of dropping out) to zero, $U_{0,i,t} - \epsilon_{0,i,t} = 0$, and include the effective price $p_{i,t}$ of purchasing a key only in the indirect utility of purchasing a key, $\text{buy}_{i,t} = 1$.

Price Endogeneity and Control Function. As mentioned above, the effective price $p_{i,t}$ could correlate with the residual error term $\epsilon_{1,i,t}$ and be endogenous. Given the player-level nature of the data, we cannot rely on the standard instrumental variable techniques to address price endogeneity typically used in demand estimation (Berry et al., 1995; Nevo, 2001a). However, we mitigate price endogeneity in two other ways. First, we include i's ability among the observed regressors X_i—thus removing from the unobservable $\epsilon_{1,i,t}$ the most problematic omitted variable. Second, we estimate the parameters of model (5) on the basis of a control function approach (Blundell and Powell, 2004; Blundell et al., 2013).

Our control function relies on an instrument $Z_{i,t}$ for price $p_{i,t}$ obtained from the randomness in the difficulty of each level across players (see Section 3.2). Controlling for i's ability, random variation in a level's difficulty prior to reaching pay-gate t will induce random variation in i's incentives to purchase and spend virtual coins on items other than a key (e.g., more lives or moves to clear the level), and consequently in $p_{i,t}$. In practice, we define $Z_{i,t}$ as the number of times we observe i being close to failing any of the levels between pay-gates $t - 20$ and t, where we consider i "being close to failing" level ℓ as i's score in ℓ within a 5% interval below the ℓ-specific score threshold necessary to clear level ℓ:[16]

$$Z_{i,t} = \sum_{\ell=1}^{t} 1\left(0.95 \times \text{Necessary Score}_\ell < \text{Score}_{i,\ell} < \text{Necessary Score}_\ell\right). \tag{7}$$

We follow Petrin and Train (2010) and implement the control function approach to estimate model (5) as follows. Given the instrument $Z_{i,t}$ in (7), we assume that $p_{i,t}$ is given by:

$$p_{i,t} = \zeta_t + X_i\gamma + \lambda Z_{i,t} + \mu_{i,t}, \tag{8}$$

where ζ_t is an intercept and $\mu_{i,t}$ is an unobserved component of effective price potentially

[16] In Appendix C.1, we discuss this instrument in more detail and report the first step estimates of equation (8) along with alternative specifications of the instrument. Overall, estimation results are robust to alternative specifications of $Z_{i,t}$.

correlated with $\epsilon_{1,i,t}$ (causing price endogeneity) but independent of $(\epsilon_{0,i,t}, \epsilon_{2,i,t})$. We also assume that the expectation of $\epsilon_{1,i,t}$ conditional on $\mu_{i,t}$ is linear:[17]

$$\epsilon_{1,i,t} = \theta\mu_{i,t} + \tilde{\epsilon}_{1,i,t}, \tag{9}$$

where $\theta\mu_{i,t}$ is our control function. Finally, by substituting (9) back into (5), defining $(\tilde{\epsilon}_{0,i,t}, \tilde{\epsilon}_{2,i,t}) = (\epsilon_{0,i,t}, \epsilon_{2,i,t})$, $V_{\text{buy},i,t}(\eta_i) = U_{\text{buy},i,t}(\eta_i) - \tilde{\epsilon}_{\text{buy},i,t}$ for each $\text{buy}_{i,t} \in \{0, 1, 2\}$, and assuming that $(\tilde{\epsilon}_{0,i,t}, \tilde{\epsilon}_{1,i,t}, \tilde{\epsilon}_{2,i,t})$ are i.i.d. Gumbel (McFadden, 1974), we obtain our mixed logit specification of discrete choice model (2):

$$Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t}) = \int \frac{\exp(V_{1,i,t}(\eta_i))}{1 + \exp(V_{1,i,t}(\eta_i)) + \exp(V_{2,i,t}(\eta_i))} \phi(\eta_i|\Sigma) d\eta_i, \tag{10}$$

where $\phi(\cdot|\Sigma)$ is the normal density of η_i in (6) with Σ denoting its variance-covariance matrix and $V_{1,i,t}(\eta_i) = \delta_1 + \delta_{1,t} + \delta_{1,i} + X_i\beta_1 - (\alpha + \alpha_t + \alpha_i + X_i\pi)p_{i,t} + \theta\mu_{i,t}$ includes control function $\theta\mu_{i,t}$, based on (7), (8), and (9), to account for the potential endogeneity of $p_{i,t}$.

Estimation Results. In the absence of experimental variation, only players with $sg_{i,t} > 0$ face a locked pay-gate t and so the choice of whether to unlock it. This raises the concern that $Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t}, sg_{i,t} > 0) \neq Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t})$, which complicates the identification of mixed logit model (10). As mentioned above, however, the experimental variation described in Section 3.2 allows us to overcome this form of endogenous selection by estimating the model on the sample of players in the No Stars group, who cannot use their accumulated stars to unlock pay-gates and always face $\text{lock}_{i,t} = 1$.

Using the sample of players in the No Stars group, we address the additional concern of price endogeneity by estimating (10) on the basis of the control function approach proposed by Petrin and Train (2010). We first estimate (8) by OLS, compute each $\hat{\mu}_{i,t}$ as the fitted residual of that regression, then plug $\hat{\mu}_{i,t}$ in $V_{1,i,t}(\eta_i)$, and finally estimate mixed logit model (10) by Simulated Maximum Likelihood using 100 random Halton sequences per player (Bhat, 2003). We compute the variance-covariance matrix of the estimator as in Karaca-Mandic and Train (2003) to account for the two-step nature of the control function procedure. We report the results for both estimation steps in Appendix C, while here we visually summarize the implied estimated price elasticities.

Figure 9 plots the distribution of the estimated price elasticities at pay-gate 40 evaluated at effective prices $p_{i,40} = 10, 30, 70, 140$ virtual coins (see Appendix D.1 for the computational details). Each panel plots the distribution of price elasticities across players when everyone faces the same effective price $p_{i,40}$.[18] Two intuitive findings emerge from Figure 9: first,

[17]We attempted the estimation of model (5) on the basis of various—more elaborate—specifications of both (8) and (9), but found no substantial differences. As a consequence, we decided to stick to these simpler and similarly effective linear specifications.

[18]Importantly, while we relied on the players in the No Stars for the estimation of model (10) to avoid

for any given effective price, there is heterogeneity across players facing the same pay-gate; second, as $p_{i,40}$ increases, a player's demand quickly becomes extremely elastic.

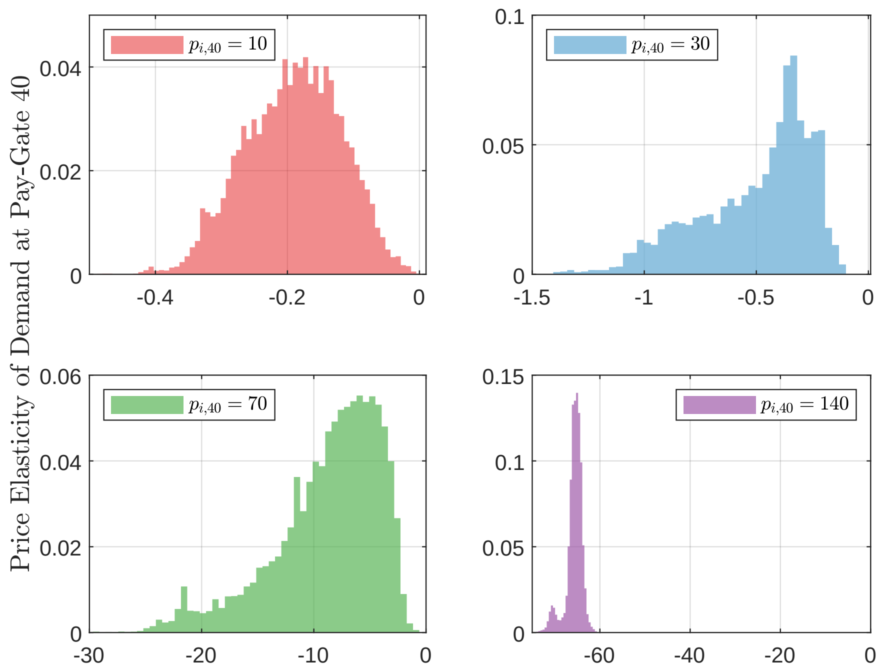

Notes: These histograms display the distribution of the price elasticity of demand at pay-gate 40 evaluated at different effective prices $p_{i,40} = 10, 30, 70, 140$. Price elasticity of demand is defined in Appendix D.1. In each panel, the distribution of the price elasticity is evaluated at the same effective price $p_{i,40}$ for all players. The sample used to compute these price elasticities includes all players of Group 40 (see footnote 19).

Figure 9: Price Elasticity of Demand at Pay-Gate 40

Figure 10 compares average price elasticities of demand across pay-gates when evaluated at a given effective price. The left panel plots results for effective prices ranging from 0 to 40 virtual coins, while the right panel plots results for effective prices ranging from 0 to 140 virtual coins (note the much larger scale on the y-axis). Figure 10 makes clear that price elasticity is heterogeneous not only across players at a given pay-gate, but also across pay-gates. This is driven by heterogeneity across players dropping out of the game and thus "surviving" at each pay-gate. Moreover, the price elasticity seems to substantially increase (i.e., become negative) for all pay-gates once the effective price goes beyond 50 virtual coins (right panel).

sample selection complications, here we only plot the implied price elasticities for the players in Group 40, who face the default design of the game. See also footnote 19.

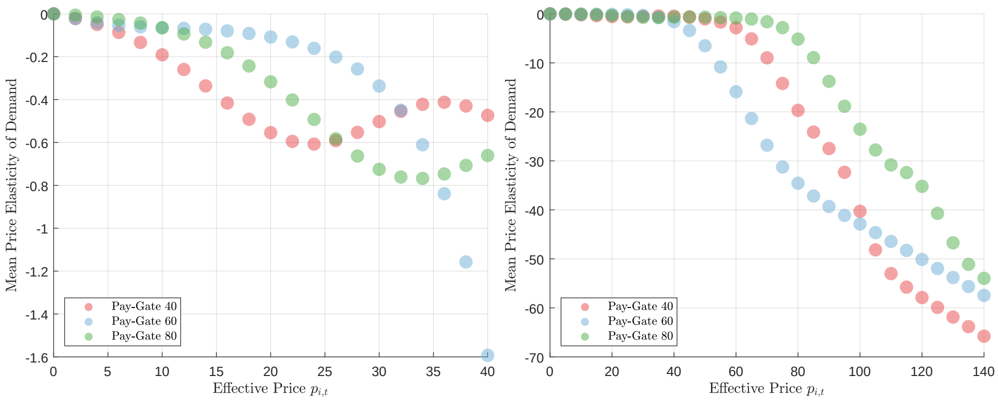

Notes: This figure displays the average price elasticity of demand across different pay-gates for given effective price. The left panel plots results for effective prices ranging from 0 to 40 virtual coins while the right panel for effective prices ranging from 0 to 140 virtual coins. For each given value of effective price p on the x-axis, we plot the average price elasticity of demand of each pay-gate so that $p_{i,40} = p_{i,60} = p_{i,80} = p$. The price elasticity of demand is defined in Appendix D.1. The sample includes all players of Group 40 (see footnote 19).

Figure 10: Average Price Elasticity of Demand across Pay-Gates

Finally, Figure 11 displays two binned scatter plots of the price elasticity of demand at pay-gate 40 with respect to a player's ability (left panel) and the log(GDP per capita) of their country (right panel). The left panel shows a striking relationship between price elasticity and player's ability. Less able players are more inelastic with respect to the effective price $p_{i,t}$. For a given increase in $p_{i,t}$, they are more likely to purchase keys using real money rather than to stop playing or ask a friend to download the game. The right panel of Figure 11 then confirms an intuitive positive relationship between price elasticity and GDP per capita of a player's country: for given increase in $p_{i,t}$, players from wealthier countries are far less price elastic at pay-gate 40 than players from poorer countries.

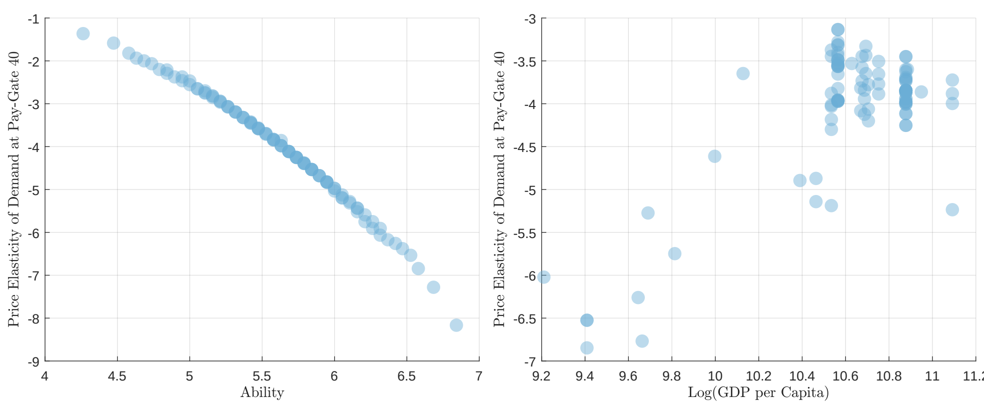

Notes: This figure shows two binned scatter plots of the price elasticity of demand at pay-gate 40 with respect to a player's ability (left panel) and the log(GDP per capita) of their country (right panel). Price elasticity of demand is defined in Appendix D.1. Players' ability and log(GDP per capita) are defined in Section 3.1. In each panel, we segment the x-axis in 100 equally sized groups. For each of these groups, we then plot the average price elasticity of demand on the y-axis. The sample includes all players of Group 40 who reached pay-gate 40 with a positive star gap (see footnote 19). The effective prices used to calculate the price elasticity of demand are based on the empirical distribution of effective prices as described in Appendix D.2.

Figure 11: Price Elasticity of Demand at Pay-Gate 40 by Ability and Log(GDP per Capita)

Overall, these results suggest that, choosing effective prices on the basis of routinely collected data may be profitable for the firm. For example, Figure 10 suggests that the firm may gain by setting effective prices somewhere around 40-50 virtual coins, higher than the observed average effective prices of around 35 virtual coins. In addition, Figure 11 confirms the importance of observing and controlling for player's ability in studying the pricing strategies of the firm, something that was not possible until a few years ago with standard offline games.

5.3 Model Validation

In Appendix E, we conduct some model validation analysis and illustrate the estimated model's ability to predict player behavior under counterfactual pricing strategies.

6 Simulation of Alternative Pricing Strategies

In this last Section, we rely on our estimated model to evaluate the returns of alternative pricing strategies for the firm. To provide intuition, we first highlight some of the most salient trade-offs faced by the firm when choosing effective prices. These trade-offs uncover the complex nature of the optimization problem and highlight the value of the empirical methods we employ. Second, we simulate alternative pricing strategies characterized by increasing discrimination and compare their implied expected revenues to those observed to be earned by the firm. While we relied on the players in Group 40 for the estimation of model (4) and on those in No Stars for the estimation of model (10), we perform all counterfactual simulations only with respect to the players in Group 40. Players in Group 40 face the default design of the game, which corresponds to our discrete choice model in equation (4).[19]

6.1 Understanding the Firm's Optimization Problem

Here we use our estimates of models (4) (probability of reaching the next gate) and (10) (choice of how to unlock a gate) and the simulation procedures detailed in Appendices D.1 and D.2 to investigate whether the per-player expected revenue of the firm is affected by dynamic considerations (i.e., if prices at different pay-gates should be set jointly or can instead be chosen independently) and by player heterogeneity (i.e., if the firm should condition prices on observed ability and/or GDP per capita). In all simulations, per-player expected revenue is averaged across the 43,660 players in Group 40 during the 15 days of our sample in 2013.

[19] Counterfactual simulations based on players from the other experimental groups would not be very informative, in that both players in Group 20 and in No Stars face game rules incompatible with those of the default design of the game embodied in model (4).

Dynamics across Pay-Gates. Although we show in Section 4.2 that players behave myopically, it is still possible for the firm's optimal pricing to involve dynamic considerations across pay-gates. As we saw in Figure 10, because different players drop out of the game at different levels, the price responsiveness of the "surviving" population changes at different pay-gates. Potentially, the firm could then increase expected revenue by influencing this selection mechanism with an appropriate choice of effective prices at different pay-gates. For instance, Figure 12 shows how per-player expected revenue (in $) from pay-gates 40, 60, and 80 changes as the firm sets different combinations of effective prices across pay-gates 40 and 60 (keeping $p_{i,80} = 70$). Each line represents the perimeter of an iso-revenue area, gathering all combinations of $(p_{i,40}, p_{i,60})$ that deliver an identical per-player expected revenue. Darker shades of blue are associated with lower per-player expected revenue.

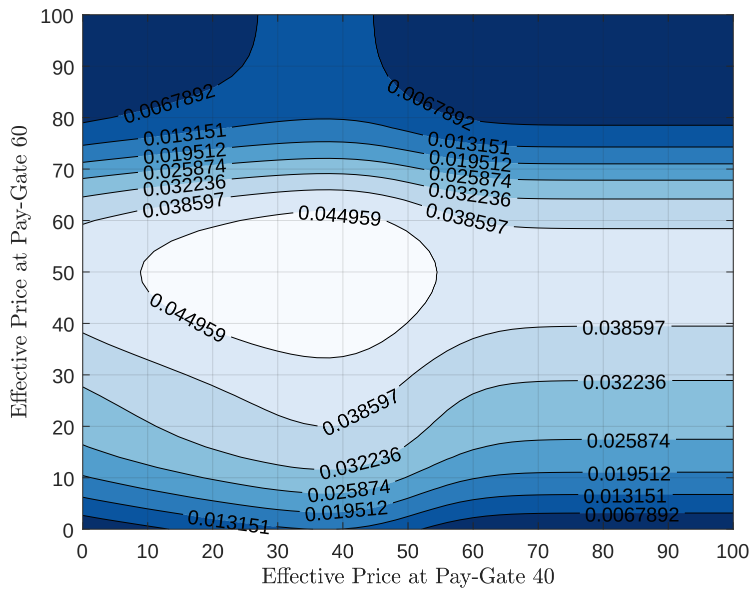

Notes: This figure compares the per-player expected revenue (in $) from pay-gates 40, 60, and 80 for different combinations of effective prices at different pay-gates. Each line represents the perimeter of an iso-revenue area, gathering all combinations of $(p_{i,40}, p_{i,60})$ (while keeping fixed $p_{i,80} = 70$) that deliver an identical per-player expected revenue. Darker shades of blue are associated to lower per-player expected revenue. The simulation of per-player expected revenue is based on our estimates of models (4) and (10) and the procedure detailed in Appendices D.1 and D.2. Per-player expected revenue is averaged across the 43,660 players in Group 40 during the 15 days of our sample in 2013.

Figure 12: Dynamic Pricing and Iso-Revenue Lines

The per-player expected revenue iso-quants depicted in Figure 12 reveal two findings. First, there is some dynamic connection between the effective prices across pay-gates. For example, for fixed $p_{i,60} = 50$ virtual coins, the firm can achieve the largest per-player expected revenue by setting $10 \leq p_{i,40} \leq 55$ virtual coins. For $10 > p_{i,40} > 55$, the firm would decrease per-player expected revenue because keys to unlock pay-gate 40 would be either too cheap or too expensive, so that too many players would not unlock pay-gate 40 and drop out of the game. Second, per-player expected revenue appears to be more responsive to changes in $p_{i,60}$ than to changes in $p_{i,40}$, stressing that not all pay-gates carry the same weight in terms of per-player expected revenue for the firm. Despite all this, it is however important to highlight that Figure 12 does not exclude the possibility that optimal effective prices across pay-gates may coincide, so that $p_{i,40} = p_{i,60}$.

The potential dynamics of the problem faced by the firm can also be directly seen by looking at the per-player expected revenue function presented in equation (3). The effective price $p_{i,t}$ plays two roles: (i) it affects i's expected revenue from pay-gate t, $Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t}) \times p_{i,t}$ and (ii) it also affects i's expected revenue from future pay-gates by changing the probability of a player dropping out, $(1 - Pr_{i,t}(\text{buy}_{i,t} = 0|\text{lock}_{i,t} = 1, p_{i,t})) \times R_{i,t+20}(p_{i,t+20}|p_{i,>t+20})$. In this sense, the effective price $p_{i,t}$ must be chosen by the firm to balance the per-player expected revenue from current pay-gate t and that from future pay-gates $t' > t$. In Figure 13, we separately illustrate these by plotting the current (i) and future (ii) components of per-player expected revenue (in \$) from pay-gate 40 as a function of $p_{i,40}$, holding fixed $(p_{i,60}, p_{i,80}) = (70, 70)$.

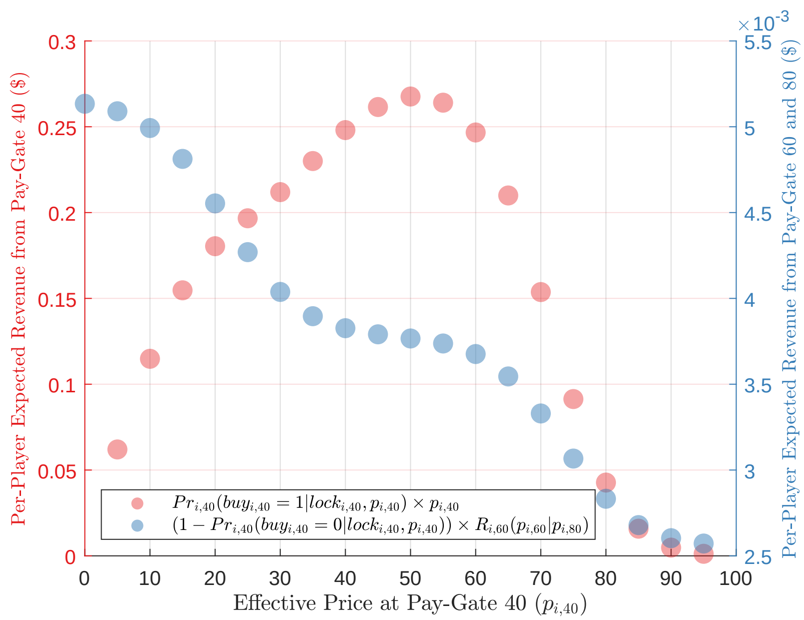

Notes: This figure displays the current and future components of per-player expected revenue (in \$) from pay-gate 40 as a function of $p_{i,40}$, as described in equation (3). On the x-axis we report different values of the effective price at pay-gate 40 ($p_{i,40}$) and on the y-axis we show: on the left, the per-player expected revenue from pay-gate 40 denoted by $Pr_{i,40}(\text{buys}_{i,t} = 1|\text{lock}_{i,40} = 1, p_{i,40}) \times p_{i,40}$); on the right, the per-player expected revenue from future pay-gates 60 and 80 evaluated at $(p_{i,60}, p_{i,80}) = (70, 70)$ and denoted by $(1 - Pr_{i,40}(\text{buys}_{i,t} = 0|\text{lock}_{i,40} = 1, p_{i,40})) \times R_{i,60}(p_{i,60} = 70|p_{i,80} = 70)$. The simulation of per-player expected revenue is based on our estimates of models (4) and (10) and the procedure detailed in Appendices D.1 and D.2. Per-player expected revenue is averaged across the 43,660 players in Group 40 during the 15 days of our sample in 2013.

Figure 13: Current and Future Components of Per-Player Expected Revenue

On the one hand, by increasing $p_{i,40}$, per-player expected revenue from *future* pay-gates 60 and 80 decreases due to a decrease in the probability of purchasing a key which is not compensated by an increased probability to ask a friend to download the game to unlock pay-gate 40, $(1 - Pr_{i,40}(\text{buy}_{i,40} = 0|\text{lock}_{i,40} = 1, p_{i,40}))$. On the other, by increasing $p_{i,40}$, per-player expected revenue from *current* pay-gate 40 increases up to $p_{i,40} = 50$ and quickly falls afterward. Importantly though, note that the scale of the y-axis on the right-hand side of Figure 13 is of an order of magnitude smaller than that on the left-hand side. This suggests that the per-player expected revenue from current pay-gates could be what really matters when choosing effective prices, and that ignoring this inter-pay-gates trade-off may not be very costly for the firm.

Player Heterogeneity. Any given price change will not impact homogeneously the expected revenue from different players. As pointed out in Figure 11, this is due to player heterogeneity in terms of ability and GDP per capita, which translates into heterogeneous price sensitivities. We explore this in Figure 14 by comparing the per-player expected revenue from pay-gates 40, 60, and 80 for various effective prices across players with different ability (left panel) and from countries with different GDP per capita (right panel). In each panel, we set effective prices across all pay-gates and players to be uniform and equal to, in turn, 10, 30, and 70. Each point represents the average simulated per-player expected revenue for the 5% of players closest to the value of ability or log(GDP per capita) on the x-axes.

By looking at any binned scatter plot of the same color, we note that any uniform effective price leads to a different per-player expected revenue depending on a player's ability or log(GDP per capita). In addition, by comparing the vertical distances among plots of different colors, we also observe that different uniform effective prices will differently impact players with heterogeneous levels of ability or log(GDP per capita). For example, while the three uniform prices considered lead to similar per-player expected revenues for players with lower ability (up to 5), more able players (with ability larger than 5) generate up to double the amount of per-player expected revenue when the uniform price is 30 as opposed to 10 or 70.

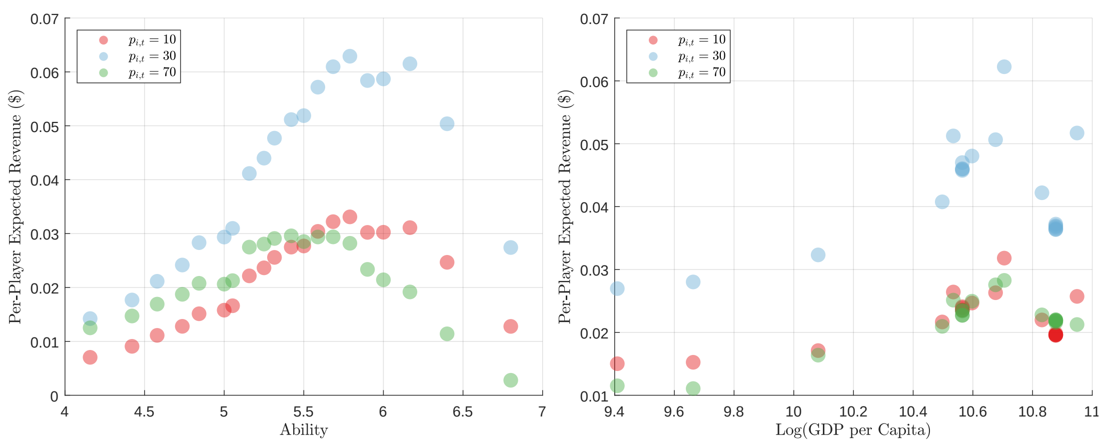

Notes: These binned scatter plots compare the per-player expected revenue from pay-gates 40, 60, and 80 for various effective prices across players with different ability (left panel) and from countries with different GDP per capita (right panel). In each panel, we set effective prices across all pay-gates and players to be uniform and equal to, in turn: 10, 30, and 70. Each point represents the average simulated per-player expected revenue for the 5% of players closest to the value of ability or log(GDP per capita) on the x-axis. Ability and log(GDP per capita) are described in Section 3.1. The simulation of per-player expected revenue is based on our estimates of models (4) and (10) and the procedure detailed in Appendices D.1 and D.2. Per-player expected revenue is averaged across the 43,660 players in Group 40 during the 15 days of our sample in 2013.

Figure 14: Per-Player Expected Revenue by Ability and GDP per Capita for Alternative Prices

6.2 Simulation Results

Next, we combine our estimates of models (4) (probability of reaching the next gate) and (10) (choice of how to unlock a gate), as well as the simulation procedures detailed in Appendices D.1 and D.2 to investigate the welfare effects of alternative pricing strategies, each requiring different levels of sophistication and amounts of data. All simulations are based on the 43,660 players in Group 40 during the 15 days of our sample in 2013.

We simulate alternative pricing strategies in which the firm directly chooses $p^*_{i,t}$ for all players and pay-gates $t = 40, 60, 80$ by maximizing per-player expected revenue from the perspective of level zero (just before players start the game) under various constraints on the flexibility of prices across players and/or pay-gates.[20] We compare the relative performance of the following pricing strategies, ordered in terms of increasing discrimination from uniform to first-degree price discrimination (see Appendix D.2 for the details):

- Observed. The observed pricing chosen by the firm, where each $p^*_{i,t}$ equals 70 virtual coins minus i's remaining endowment when facing pay-gate t.
- Uniform (70). All players face the same effective price $p^*_{i,t} = 70$. This amounts to setting everybody's endowment of virtual coins to zero from the beginning of the game.
- Uniform (Optimal). The firm optimally chooses either only one effective price p^* for all players and pay-gates (static) or an effective price p^*_t common across players but specific to each pay-gate t (dynamic).
- GDP per Capita. Third-degree price discrimination based on the observed GDP per capita of a player's country. As for "Uniform (Optimal)," we consider both a static version in which the effective prices are identical across pay-gates and a dynamic version in which the effective prices are also allowed to change across pay-gates.
- Ability. Third-degree price discrimination based on the observed gaming ability of each player.[21] We again consider both a static (identical effective prices across pay-gates) and a dynamic version (potentially different effective prices across pay-gates).
- Individual Level. First-degree price discrimination where the firm is free to choose a different effective price for each player. We again consider both a static (identical effective prices

[20] The fact that in our counterfactuals the firm *directly* chooses effective prices corresponds to restricting players' freedom to use the initial endowment of 70 virtual coins. Independently of i's endowment at pay-gate t, to progress in the game, i must use real money to purchase a key at the effective price of $p_{i,t}$ virtual coins. This greatly simplifies our model and simulations because we can proceed without specifying and estimating a further choice model for the allocation of the initial endowment of virtual coins. We believe this assumption is without loss of generality: the firm could always change any feature of the initial endowment, such as restricting the way players are allowed to use it, changing its magnitude (allowing players to have more or less than 70 virtual coins), or even removing it altogether (every player gets an endowment of zero virtual coins).

[21] As described in Section 3.1, we compute ability from each player's performance during the first 20 rounds of the game, something observed by the firm by the time the player reaches the first pay-gate at level 40.

Pricing Strategy	Static Pricing					Dynamic Pricing				
	Effective Price		Per-Player Revenue ($)			Effective Price		Per-Player Revenue ($)		
	mean	s.d.	mean	s.d.	%	mean	s.d.	mean	s.d.	%
Observed	35.566	34.529	0.011	0.108	-	-	-	-	-	-
Uniform (70)	70.000	-	0.022	0.012	93.8%	-	-	-	-	-
Uniform (Optimal)	45.000	-	0.049	0.024	340.0%	10.000	2.449	0.051	0.026	358.9%
GDP per Capita	44.500	1.500	0.049	0.025	340.7%	51.167	11.950	0.051	0.026	359.7%
Ability	44.500	3.841	0.049	0.025	343.7%	52.833	12.429	0.051	0.026	362.8%
Individual Level	45.166	5.065	0.050	0.025	346.9%	52.617	12.756	0.052	0.026	368.1%

Notes: This table summarizes our counterfactual simulation results in terms of effective prices and per-player expected revenues. Each row refers to a pricing strategy and summarizes the simulated effective prices chosen by the firm (in virtual coins, where $1 ≈ 70 virtual coins) and the corresponding per-player expected revenues (in $). The columns denoted by "%" report the percentage increase in per-player expected revenue implied by the row pricing strategy with respect to the observed pricing chosen by the firm (i.e., 0% means same average as the observed pricing). All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. The left panel summarizes results for the case in which effective prices do not change among pay-gates (static pricing). The right panel instead summarizes results for the case in which effective prices are allowed to change also among pay-gates (dynamic pricing). All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Table 7: Counterfactual Pricing Strategies, Effective Prices and Expected Revenues

Pricing Strategy	Static Pricing				Dynamic Pricing			
	Δ Consumer Surplus ($)		Δ Total Surplus ($)		Δ Consumer Surplus ($)		Δ Total Surplus ($)	
	mean	s.d.	mean	s.d.	mean	s.d.	mean	s.d.
Uniform (70)	-0.0187	0.0137	-0.0083	0.0121	-	-	-	-
Uniform (Optimal)	-0.0369	0.0211	0.0008	0.0079	-0.0346	0.0203	0.0052	0.0097
GDP per Capita	-0.0376	0.0223	0.0002	0.0090	-0.0344	0.0211	0.0055	0.0105
Ability	-0.0383	0.0227	-0.0002	0.0069	-0.0334	0.0204	0.0068	0.0086
Individual Level	-0.0380	0.0241	0.0005	0.0089	-0.0336	0.0224	0.0072	0.0108

Notes: This table summarizes our counterfactual simulation results in terms of per-player consumer surplus and per-player total surplus, computed as the sum between changes in per-player expected revenue and in per-player consumer surplus. Each row refers to a pricing strategy and summarizes the simulated change in per-player consumer surplus and in per-player total surplus (both in $) with respect to the observed pricing. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. The left panel summarizes results for the case in which effective prices do not change among pay-gates (static pricing). The right panel instead summarizes results for the case in which effective prices are allowed to change also among pay-gates (dynamic pricing). All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Table 8: Counterfactual Pricing Strategies, Consumer Surplus and Total Surplus

across pay-gates) and a dynamic version (potentially different effective prices across pay-gates).

Tables 7 and 8 summarize our counterfactual simulation results. Table 7 reports results of counterfactual effective prices and expected revenues. Table 8 reports changes in consumer surplus and total surplus, computed as the sum between changes in expected revenues and in consumer surplus. Each row of Table 7 refers to a pricing strategy and reports mean and standard deviation of the simulated effective prices chosen by the firm (in virtual coins) and of the corresponding per-player expected revenues (in $). The columns denoted by "%" report the percentage increase in per-player expected revenue implied by the row pricing strategy with respect to the observed pricing chosen by the firm (0% means same average as the observed pricing). Analogously, each row of Table 8 refers to a pricing strategy and summarizes the simulated change in per-player consumer surplus and in per-player total surplus (both in $) with respect to the observed pricing. Figures 15–18 visualize these results by plotting the simulated distributions of effective prices, per-player expected revenue, changes in per-player consumer surplus, and changes in per-player total surplus for the "static" pricing strategies considered in the left panels of Tables 7 and 8. Appendix Figures A-13, A-14, A-18, and A-23 plot analogous simulated distributions for the "dynamic" pricing strategies considered in the right panels of Tables 7 and 8.

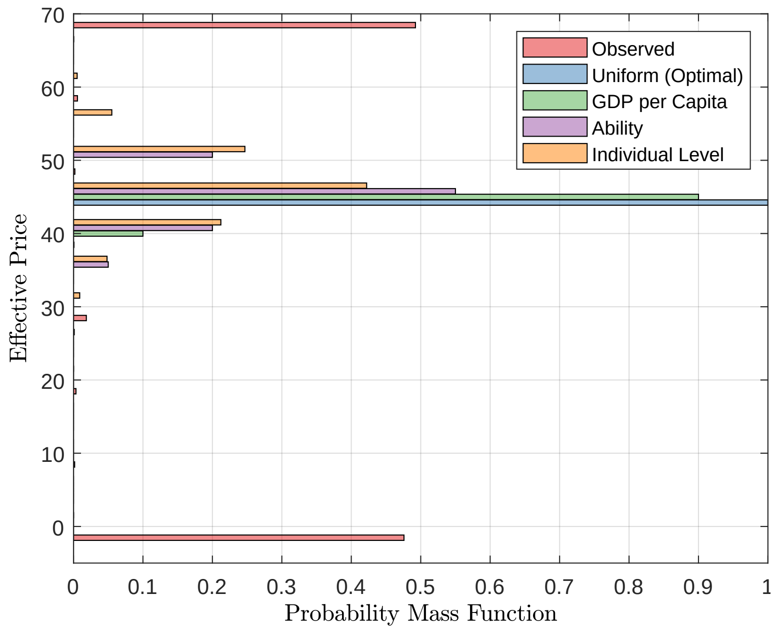

Notes: This figure shows the simulated distribution of effective prices (in virtual coins, where \$1 ≈ 70 virtual coins) across players and pay-gates for the "static" pricing strategies considered in the left panel of Table 7. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure 15: Distribution of Effective Prices in Static Pricing Strategies

Our counterfactual simulation results accord with intuition: more flexible and discriminatory pricing strategies lead to larger per-player expected revenue at the expense of lower per-player consumer surplus. Uniform pricing is associated with lower per-player expected revenue than discriminatory pricing strategies. Similarly, dynamic pricing strategies that allow effective prices to vary across pay-gates lead to higher per-player expected revenue than their more restrictive static counterparts. While it is well known that price discrimination will in general enable a monopolist to seize larger portions of consumer surplus and thus increase profit at the expense of consumer surplus (Varian, 1989), Tables 7 and 8 provide three striking and perhaps less obvious insights.

First, by comparing the first three rows of Table 7, from the firm's perspective the free endowment of 70 virtual coins given to players when they begin playing is too large. The resulting average observed effective price of 35.6 virtual coins is too low. If the firm were to remove this free endowment and charge a uniform effective price of 70 virtual coins (as in "Uniform (70)," second row of Table 7), per-player expected revenue would almost double (+93.8%). In addition, if it were also to optimally adjust the uniform effective price to 45 virtual coins (as in "Uniform (Optimal)," third row of Table 7), per-player expected revenue would more than quadruple (+340%). Observed pricing seems far from maximizing profit and the firm could substantially be better-off by limiting itself to the use of uniform pricing. A possible explanation is that the game was only launched three months before our data were collected, while the firm was probably still learning how to profitably set effective prices (Dube and Misra, 2019; Huang et al., 2020). The fact that the firm may behave sub-optimally does not affect the validity of our counterfactual simulations. Indeed, because the marginal cost of each additional player is zero, we only rely on the estimation of demand for game content (which does not require any profit-maximization assumption).[22]

Second, while more flexible and discriminatory pricing strategies lead to larger per-player expected revenue, the relative gains from their implementation are limited when compared to a simple uniform pricing strategy. This can be seen by comparing the third with the last three rows (fourth to sixth) of Table 7. Despite the different distributions of effective prices implied by each pricing strategy (Figure 15), the corresponding distributions of per-player expected revenues are remarkably similar (Figure 16), with the relative gains of static "Individual Level" ($45,2$ of average effective price) with respect to static "Uniform (Optimal)" (a unique effective price) being essentially negligible ($\$0.050 - \$0.049 = \$0.001$). Appendix Figure A-15 further stresses this point: we construct 20 groups of players based on players' ability (left panel) and GDP per capita (right panel) and plot the average group-specific difference in per-player expected revenue. Remarkably, "Uniform (Optimal)" performs almost as well as any discriminatory pricing strategy not only on average (left panel of Table 7) but also conditional on player's ability and GDP per capita (Appendix Figure A-15).

The right panel of Table 7 and Appendix Figures A-13, A-14, and A-16 tell a similar story also for dynamic versions of these pricing strategies, underlying that even when the effective prices can change across pay-gates, "Uniform (Optimal)" (three prices $p^*_{40}, p^*_{60}, p^*_{80}$) still seizes most of the potential revenue of "Individual Level" (three prices $p^*_{i,40}, p^*_{i,60}, p^*_{i,80}$ for each $i = 1, ..., 43660$). Appendix Figure A-17 then compares the relative gains of implementing a dynamic versus a static version of each pricing strategy by ability group and GDP per capita. Consistent with the findings that players behave myopically and that inter-pay-gates trade-offs may not be very relevant for the firm (Section 4.2 and Figures 12 and 13), Appendix Figure A-17 shows that while dynamic pricing strategies slightly outperform their static counterparts, the implied relative gains are in practice very limited (note the smaller order of magnitude of the y-axis with respect to Appendix Figures A-15 and A-16).

[22]Deviations from profit-maximization instead represent a problem when the simulation of counterfactuals also requires the estimation of marginal cost functions, which typically hinges on the correct specification of the optimization problem solved by the firm (Berry et al., 1995; Nevo, 2001b).

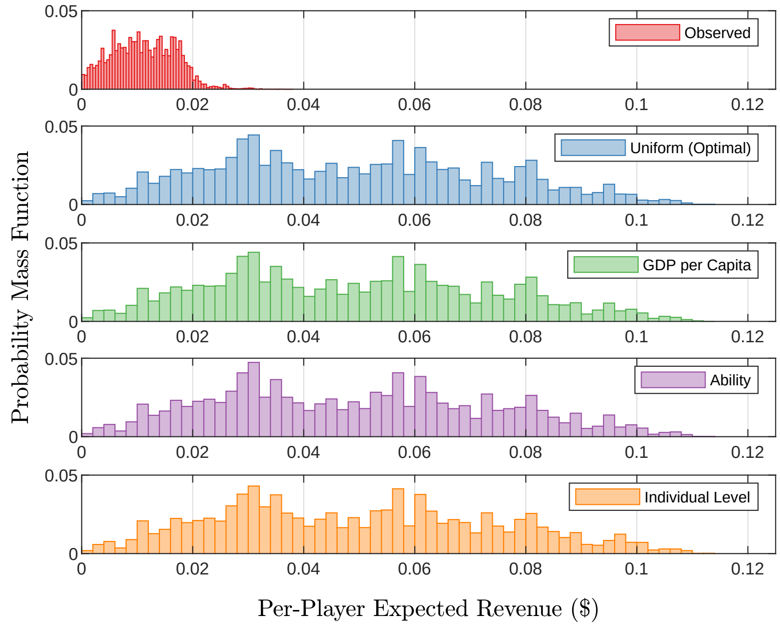

Notes: This figure shows the simulated distribution of per-player expected revenue (in $) across players for the "static" pricing strategies considered in the left panel of Table 7. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure 16: Distribution of Per-Player Expected Revenue, Static Pricing Strategies

Mirroring these findings, Table 8 illustrates that each of the alternative pricing strategies would lead to a loss in per-player consumer surplus. This can be clearly seen in Figure 17, which shows that the distribution of changes in per-player consumer surplus associated to each alternative pricing strategy would always have negative support. This makes intuitive sense, in that each alternative pricing strategy would imply a higher average effective price than the observed one (of around 10 virtual coins, see left panel of Table 7), enabling the firm to extract more of the players' surplus. Importantly, mixed logit model (10) allows players to drop out of the game at any pay-gate t (by choosing $\text{buy}_{i,t} = 0$) if, for example, effective prices were "too high." In other words, this simulated extraction of consumer surplus is not conditional on the players being held "captive" in the game, but it is rather based on a more effective exploitation of their preferences. Appendix Figures A-18– A-22 visualize additional dimensions of heterogeneity. Although also dynamic counterfactual pricing strategies induce losses in consumer surplus (Appendix Figure A-18), they usually generate smaller decreases than their static counterparts (Appendix Figure A-19).

Third, by summing the increases in expected revenue and the decreases in consumer surplus, Table 8 shows that, on average, the per-player total surplus implied by most of the alternative pricing strategies would be non-negative (with the exception of "Uniform (70)" and of the static version of "Ability"). On average, these counterfactual pricing strategies would generate enough additional expected revenue to compensate the corresponding loss in consumer surplus. As the right panel of Table 8 illustrates, the dynamic pricing strategies would perform slightly better than their static counterparts as a result of a slightly larger increase in expected revenue (Appendix Figure A-17) and a slightly smaller decrease in consumer surplus (Appendix Figure A-19). The fact that average per-player total surplus is non-negative suggests that these pricing strategies would not only enable the firm to extract more of the players' surplus, but that they would also not lead to sizeable dead-weight losses—despite the increase in average effective price of around 10 virtual coins. Figure 18 and Appendix Figure A-23 highlight the distributional content of this result, stressing that—despite the non-negative average—, there would always be groups of players associated to negative changes in total surplus.

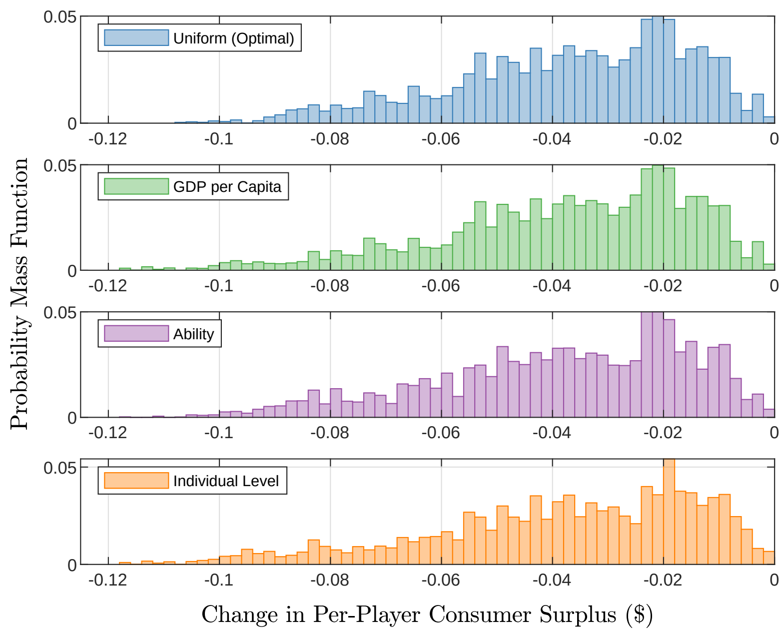

Notes: This figure shows the simulated distribution of changes in per-player consumer surplus (in $) across players for the "static" pricing strategies considered in the left panel of Table 7 as opposed to the observed pricing. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure 17: Distribution of Δ Per-Player Consumer Surplus, Static versus Observed Pricing

To summarize, our counterfactual simulation results suggest that (i) observed pricing is far from profit maximizing and leaves a lot of surplus in hands of players, (ii) optimal uniform pricing would generate most of the returns associated with more complex pricing strategies, and (iii) each of the pricing strategies considered—including optimal uniform pricing—would induce a transfer of surplus from the players to the firm without, however, generating any sizeable dead-weight loss on average.

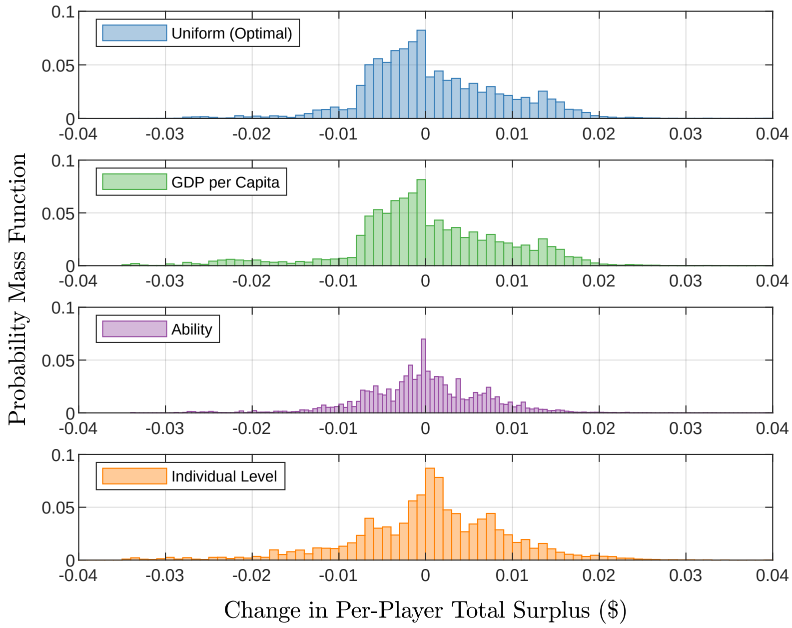

Notes: This figure shows the simulated distribution of changes in per-player total surplus (in $) across players for the "static" pricing strategies considered in the left panel of Table 8 as opposed to the observed pricing. Changes in per-player total surplus are computed as the sum between changes in per-player expected revenues and in per-player consumer surplus. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure 18: Distribution of Δ Per-Player Total Surplus, Static versus Observed Pricing

Findings consistent with (i) have been documented in other industries: because of path-dependence, imperfect information, learning, or conflicting incentives, sometimes for-profit firms do not maximize profit (Cho and Rust, 2010; DellaVigna and Gentzkow, 2019; Dube and Misra, 2019; Fioretti, 2020; Hortaçsu et al., 2021; Huang et al., 2020; Orbach and Einav, 2007). Finding (ii) is in line with Chu et al. (2011), who show in the context of a theater company that simple pricing rules can sometimes generate almost as much profit as complex ones that would however be hard to implement. Finding (ii) is also close in spirit to Levitt et al. (2016), who document limited gains of second-degree price discrimination for a large online gaming firm, and more in general to the empirical literature on the trade-offs of price discrimination and personalized pricing in the era of big data (Rossi et al., 1996; Shiller and Waldfogel, 2011; Shiller, 2015; Waldfogel, 2015). Limited gains from price discrimination may partly explain why it is rarely observed in business practice, where additional risks tied to consumer backlash and regulatory scrutiny also need to be considered (Council of Economic Advisors, 2015; DellaVigna and Gentzkow, 2019).

In contrast to our results, however, Dube and Misra (2019) document substantial returns of personalized pricing for a digital recruiting firm, highlighting the need for caution in drawing general conclusions. While we do not find any such evidence, in other digital contexts more complex pricing strategies may be much more profitable. That being said, both our results and Dube and Misra (2019) stress the large potential of "empirical" pricing rules. In our context, the firm could increase per-player expected revenues more than fourfold by optimally choosing a uniform effective price on the basis of detailed data and appropriate empirical methods. Importantly, finding (iii) stresses that, although these increases in profit would necessarily come at the expense of decreases in consumer surplus, the pricing strategies considered would not generate average losses in total welfare.

6.3 Robustness Checks

In Appendix F.1, we repeat all counterfactual simulations accounting for the predictive biases of our estimated model as documented in Appendix E. We do this by limiting our counterfactual simulations to the sub-sample of players in Group 40 for which the estimated model has the best predictive power. These checks show no qualitative difference in our results and suggest that theses predictive biases do not play a crucial role for our simulations exercises.

7 Conclusion

Our results indicate that the game developer can substantially increase profit by using basic information it has readily available on player characteristics and in-game behavior. As expected, the increase in profit largely results from a transfer of surplus from players to the firm (Varian, 1989). However, most of the pricing strategies considered do not decrease total surplus on average. Our results also show that a simple uniform pricing strategy may already guarantee most of the profit implied by elaborate forms of price discrimination (Chu et al., 2011), which might help explain why price discrimination has been used sparsely in online markets.

Our study analyzes price setting in a popular mobile game that, during the period of data collection, had a number of specific features. While these features and some of our modelling choices facilitate our empirical analysis, they may also limit the generality of our findings. First, no advertisement was shown in the game during data collection. With advertisement, the problem of the firm would differ, in that it could decide to trade-off revenue from in-app purchases in favor of consumption of game content, possibly by reducing the prices for in-app purchases. Second, our data show no evidence of a trade-off for the firm between revenue from pay-gates and other in-app purchases. Similar to advertisement, in other freemium apps this trade-off may be more prominent and lead to a more complex maximization problem for the firm. Third, similar to Dube and Misra (2019), our counterfactual simulations treat the firm as a monopolist and this may cause an overestimation of its market power when choosing prices for premium content. We believe this assumption to be appropriate in freemium games such as the one we study, where competition among mobile games occurs mostly before players download the game (for free) and then substantially softens after a player has downloaded and started to play the game (when the prices for premium content are incurred). However, this may be less applicable to non-freemium contexts in which competing firms charge positive prices already for the download of their apps.

References

Adams, B. and K. R. Williams (2019). Zone pricing in retail oligopoly. *American Economic Journal: Microeconomics 11*(1), 124–56.

Altman, N. S. (1992). An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician 46*(3), 175–185.

Aryal, G. and M. F. Gabrielli (2020). An empirical analysis of competitive nonlinear pricing. *International Journal of Industrial Organization 68*, 102538.

Assad, S., R. Clark, D. Ershov, and L. Xu (2020). Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841–890.

Bhat, C. R. (2003). Simulation estimation of mixed discrete choice models using randomized and scrambled halton sequences. *Transportation Research Part B: Methodological 37*(9), 837–855.

Blundell, R., D. Kristensen, and R. L. Matzkin (2013). Control functions and simultaneous equations methods. *American Economic Review 103*(3), 563–69.

Blundell, R. W. and J. L. Powell (2004). Endogeneity in semiparametric binary response models. *The Review of Economic Studies 71*(3), 655–679.

Bresnahan, T. F., J. P. Davis, and P.-L. Yin (2015). *Economic Value Creation in Mobile Applications*. University of Chicago Press.

Brown, Z. Y. and A. MacKay (2021). Competition in pricing algorithms. Technical report, National Bureau of Economic Research.

Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review 110*(10), 3267–97.

Carare, O. (2012). The impact of bestseller rank on demand: Evidence from the app market. *International Economic Review 53*(3), 717–742.

Cho, S. and J. Rust (2010). The flat rental puzzle. *The Review of Economic Studies 77*(2), 560–594.

Chu, C. S., L. Phillip, and A. Sorensen (2011). Bundle-size pricing as an approximation to mixed bundling. *American Economic Review 101*, 263–303.

Cohen, P., R. Hahn, J. Hall, S. Levitt, and R. Metcalfe (2016). Using big data to estimate consumer surplus: The case of uber. Technical report, National Bureau of Economic Research.

Council of Economic Advisors (2015). Big data and differential pricing – executive office of the president of the united states, council of economic advisors. Technical report, Executive Office of the President of the United States, Washington.

Crawford, G. S. and M. Shum (2007). Monopoly quality degradation in the cable television industry. *The Journal of Law and Economics 50*(1), 181–219.

Debeauvais, T. and C. V. Lopes (2015). Gate Me If You Can: The Impact of Gating Mechanics on Retention and Revenues in Jelly Splash. *in Proc Int. Foundations of Digital Games Conf. (J. P. Zagal, E. MacCallumStewart, and J. Togelius, eds.), Society for the Advancement of the Science of Digital Games*.

DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in us retail chains. *Quarterly Journal of Economics 134*(4), 2011–2084.

Dube, J.-P. and S. Misra (2019). Personalized pricing and customer welfare. Working paper.

Dubois, P., R. Griffith, and M. O'Connell (2017). The Effects of Banning Advertising in Junk Food Markets. *The Review of Economic Studies 85*(1), 396–436.

Einav, L. and J. Levin (2010). Empirical industrial organization: A progress report. *Journal of Economic Perspectives 24*(2), 145–62.

Ershov, D. (2018). Competing with superstars in the mobile app market. *Available at SSRN 3265662*.

Fioretti, M. (2020). Caring or pretending to care? social impact, firmsâ objectives and welfare. Technical report, Tech. rep., Mimeo, Sciences Po.

Fudenberg, D. and M. J. Villas-Boas (2006). Behavior based price discrimination and customer recognition. In T. J. Hendershott (Ed.), *Handbook on Economics and Information Systems*, pp. 377–436.

Fudenberg, D. and M. J. Villas-Boas (2012). Price discrimination in the digital economy. In M. Peitz and J. Waldfogel (Eds.), *The Oxford Handbook of the Digital Economy*.

Gandhi, A. and A. Nevo (2021). Empirical models of demand and supply in differentiated products industries. Technical report, National Bureau of Economic Research.

Garbarino, E. and S. Maxwell (2010). Consumer response to norm-breaking pricing events in e-commerce. *Journal of Business Research 63*(9), 1066 – 1072. Advances in Internet Consumer Behavior & Marketing Strategy.

Ghose, A. and S. P. Han (2014a). Estimating demand for mobile applications in the new economy. *Management Science 60*(6), 1470–1488.

Ghose, A. and S. P. Han (2014b). Estimating demand for mobile applications in the new economy. *Management Science 60*(6), 1470–1488.

Goldfarb, A. and C. Tucker (2019, March). Digital Economics. *Journal of Economic Literature 57*(1), 3–43.

Hortaçsu, A., O. R. Natan, H. Parsley, T. Schwieg, and K. R. Williams (2021). Organizational structure and pricing: Evidence from a large us airline.

Huang, Y., P. B. Ellickson, and M. J. Lovett (2020). Learning to set prices. *Available at SSRN 3267701*.

Iaria, A. and A. Wang (2021). An empirical model of quantity discounts with large choice sets.

Jacobi, L. and M. Sovinsky (2016, August). Marijuana on main street? estimating demand in markets with limited access. *American Economic Review 106*(8), 2009–45.

Karaca-Mandic, P. and K. Train (2003). Standard error correction in two-stage estimation with nested samples. *The Econometrics Journal 6*(2), 401–407.

Levitt, S. D. and J. A. List (2009). Field experiments in economics: The past, the present, and the future. *European Economic Review 53*(1), 1–18.

Levitt, S. D., J. A. List, S. Neckermann, and D. Nelson (2016). Quantity discounts on a virtual good: The results of a massive pricing experiment at king digital entertainment. *Proceedings of the National Academy of Sciences 113*(27), 7323–7328.

Li, K. J. and S. Jain (2016). Behavior-based pricing: An analysis of the impact of peer-induced fairness. *Management Science 62*(9), 2705–2721.

List, J. A. (2004). The nature and extent of discrimination in the marketplace: Evidence from the field. *The Quarterly Journal of Economics 119*(1), 49–89.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. *P. Zarembka (ed.), FRONTIERS IN ECONOMETRICS. Academic Press: New York*, 105–142.

McFadden, D. and K. Train (2000). Mixed mnl models for discrete response. *Journal of applied Econometrics 15*(5), 447–470.

McManus, B. (2007, Summer). Nonlinear pricing in an oligopoly market: The case of specialty coffee. *Rand Journal of Economics 38*, 512–532.

Miklós-Thal, J. and C. Tucker (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science 65*(4), 1552–1561.

Mitchell, T. M. (1997). *Machine Learning* (1 ed.). USA: McGraw-Hill, Inc.

Mullin, M. and R. Sukthankar (2000). Complete cross-validation for nearest neighbor classifiers. In *Proceedings of the Seventeenth International Conference on Machine Learning*, ICML '00, San Francisco, CA, USA, pp. 639â646. Morgan Kaufmann Publishers Inc.

Nevo, A. (2001a). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307–342.

Nevo, A. (2001b). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307–342.

Orbach, B. Y. and L. Einav (2007). Uniform prices for differentiated goods: The case of the movie-theater industry. *International Review of Law and Economics 27*(2), 129–153.

Petrin, A. and K. Train (2010). A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research 47*(1), 3–13.

Rosencrance, L. (2000 (accessed October 3, 2019)). *Consumers balk at variable DVD pricing*. ComputerWorld, 34 (2000), p. 4. https://www.computerworld.com/article/2597065/customers-balk-at---variable-dvd-pricing.html.

Rossi, P. E., R. E. McCulloch, and G. M. Allenby (1996). The value of purchase history data in target marketing. *Marketing Science 15*(4), 321–340.

Shiller, B. and J. Waldfogel (2011). Music for a song: An empirical look at uniform pricing and its alternatives. *The Journal of Industrial Economics 59*(4), 630–660.

Shiller, B. R. (2015). First degree price discrimination using big data. *Working Paper*.

Statista (2021). Digital media report 2021 – video games. Technical report, Statista.

Tirole, J. (1988). *The Theory of Industrial Organisation*. The MIT Press.

Todd, P. E. and K. I. Wolpin (2020). The best of both worlds: Combining rcts with structural modeling. *Journal of Economic Literature*.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

Varian, H. (1989). Price discrimination. In R. Schmalensee and R. Willig (Eds.), *Handbook of Industrial Organization* (1 ed.), Volume 1, Chapter 10, pp. 597–654. Elsevier.

Wagner, S. and J. Runge (2018, September/October). Nothing is free: Data-driven optimisation unlocks freemium business models' real potential. *European Business Review*, 47–50.

Waldfogel, J. (2015). Price discrimination goes to school. *The Journal of Industrial Economics 63*(4), 569–597.

Wen, W. and F. Zhu (2019). Threat of platform-owner entry and complementor responses: Evidence from the mobile app market. *Strategic Management Journal 40*(9), 1336–1367.

Yi, J., Y. Lee, and S.-H. Kim (2019). Determinants of growth and decline in mobile game diffusion. *Journal of Business Research 99*, 363–372.

Yin, P.-L., J. P. Davis, and Y. Muzyrya (2014). Entrepreneurial innovation: Killer apps in the iphone ecosystem. *American Economic Review 104*(5), 255–59.

Yuan, H. (2020). Competing for time: A study of mobile applications. Technical report, Working Paper.

APPENDIX

For Online Publication

Discrimination and Big Data: Evidence from a Mobile Puzzle Game

A Appendix: Assumptions

A.1 Further Evidence in Support of Assumption 1

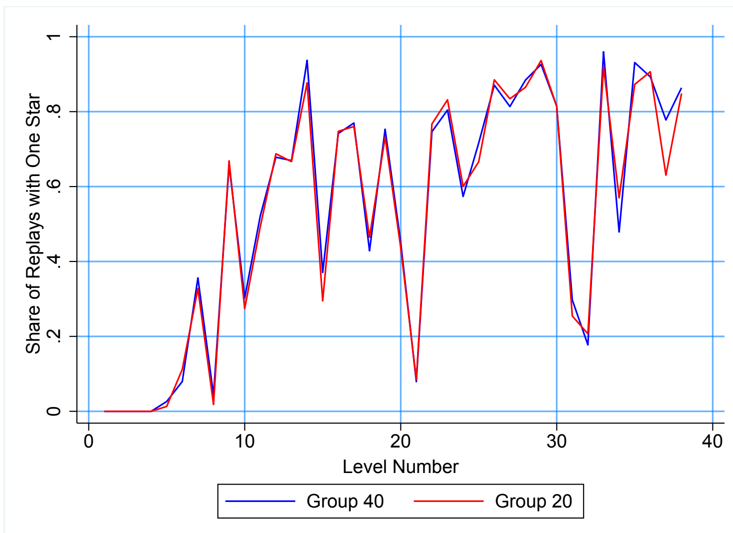

Notes: This figure compares the propensity to replay a level according to the number of available stars and according to being in Group 40 or 20. The y-axis variable is, for a given level ℓ, the share of replays which were done when only a single star had previously been collected at this level (i.e, there are two remaining stars). The sample includes players who have crossed pay-gate 20 but have no gone beyond pay-gate 40. Among these, we keep only players of Group 40 and 20 who hit pay-gate 20 with a non-positive star gap. The definition for stars is provided in Section 3.1.

Figure A-1: Group 20 vs Group 40: Prob. to re-play level initially cleared with one star

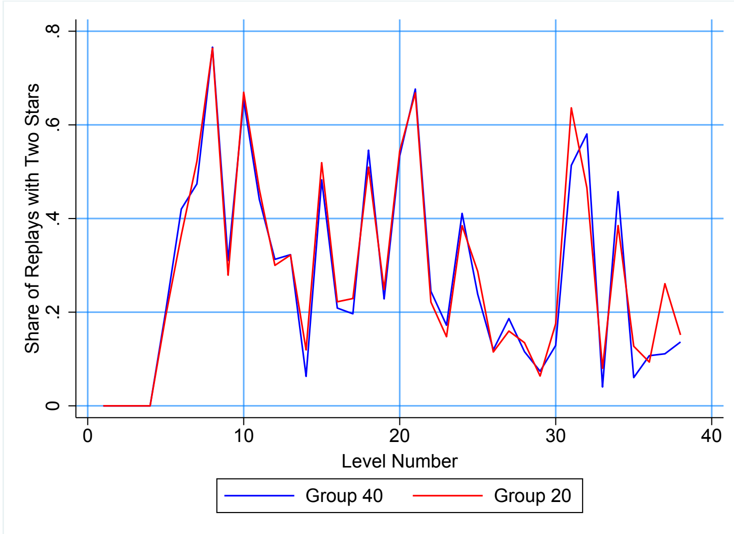

Notes: This figure compares the propensity to replay a level according to the number of available stars and according to being in Group 40 or 20. The y-axis variable is, for a given level ℓ, the share of replays which were done when two stars had previously been collected at this level (i.e, there are one remaining stars). The sample includes players who have crossed pay-gate 20 but have no gone beyond pay-gate 40. Among these, we keep only players of Group 40 and 20 who hit pay-gate 20 with a non-positive star gap. The definition for stars is provided in Section 3.1.

Figure A-2: Group 20 vs Group 40: Prob. to re-play level initially cleared with two stars

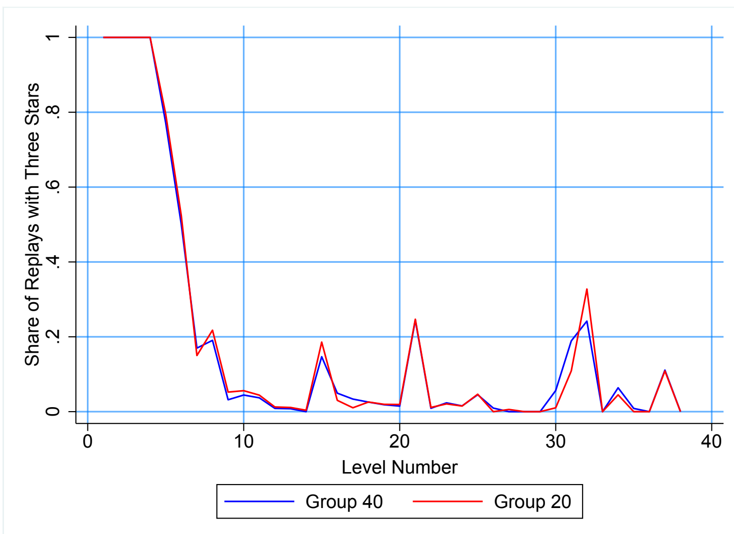

Notes: This figure compares the propensity to replay a level according to the number of available stars and according to being in Group 40 or 20. The y-axis variable is, for a given level ℓ, the share of replays which were done when three stars had previously been collected at this level (i.e, there are zero remaining stars). The sample includes players who have crossed pay-gate 20 but have no gone beyond pay-gate 40. Among these, we keep only players of Group 40 and 20 who hit pay-gate 20 with a non-positive star gap. The definition for stars is provided in Section 3.1.

Figure A-3: Group 20 vs Group 40: Prob. to re-play level initially cleared with three stars

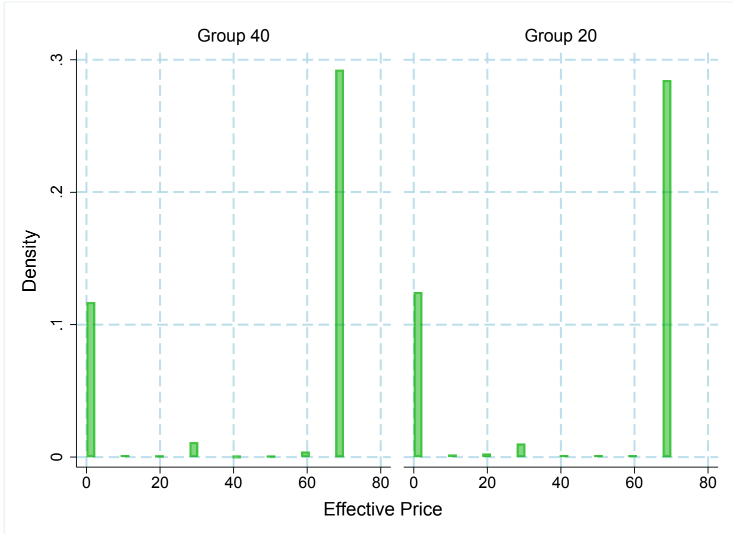

Notes: This histogram displays the effective price distribution at pay-gate 60 for Group 20 and Group 40. Each observation is a player who faces pay-gate 60 with a positive star gap. The sample includes all players of Group 40 and 20 who satisfy this condition.

Figure A-4: Distribution of effective prices $p_{i,t}$ at pay-gate $t = 60$

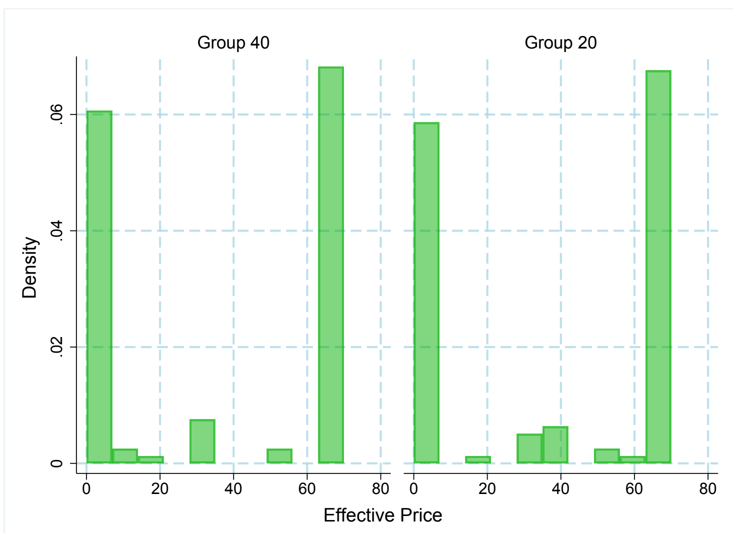

Notes: This histogram displays the effective price distribution at pay-gate 80 for Group 20 and Group 40. Each observation is a player who faces pay-gate 80 with a positive star gap. The sample includes all players of Group 40 and 20 who satisfy this condition.

Figure A-5: Distribution of effective prices $p_{i,t}$ at pay-gate $t = 80$

A.2 Relationship between Pay-Gate Purchases and Non-Pay-Gate Purchases

This Appendix provides evidence that purchases outside of pay-gates (i.e., non-pay-gate purchases within the game) are not affected by purchases at pay-gates (i.e., purchases of keys to unlock pay-gates). This is important to justify our focus on the firm's revenue from purchases at pay-gates (see Section 4.1). The evidence presented in this appendix suggests that purchases at pay-gates do not crowd out non-pay-gate purchases. This allows us to analyze revenue from purchases at pay-gates separately from other in-game purchases.

We do this by relying on the experimental design described in Section 3.2.1 and used to test Assumption 1 in Section 4.2. In particular, we compare non-pay-gate purchases by players in Group 20 who faced an additional pay-gate at level 20 with those of players in the other experimental groups who did not. We consider players in Group 20 with a positive star gap at pay-gate $t = 20$ to be "treated" with an additional pay-gate. Players of Group 40 and No Star with a positive star gap at pay-gate $t = 20$ are instead considered as the control group, because they did not face any pay-gate $t = 20$. We implement this test using two alternative measures of payment outside of pay-gates. We define two alternative variables measuring the non-pay-gate purchases by players: "Accumulated Purchases" as the total number of purchases made by a player between level $\ell = 21$ and level $\ell = 39$ and "Indicator of Purchase" as a dummy equal to one if "Accumulated Purchases" is greater than zero.

We start by comparing non-pay-gate purchasing behaviors on the basis of t-tests. Table A-1 compares the purchasing behavior of players in Group 20 versus players in No Star, while Table A-2 compares Group 20 with Group 40. In both cases and for both variables, we find no statistically significant differences.

	(1) No Star		(2) Group 20		(3) No Star−Group 20	
	mean	SE	mean	SE	diff.	t-test
Accumulated Purchases	0.061	0.671	0.053	0.506	-0.008	-1.207
Indicator of Purchase	0.023	0.151	0.023	0.149	-0.001	-0.352
Observations	39102		7409		46511	

Notes: This table presents evidence on whether facing an additional pay-gate affects the future non-pay-gate purchasing behavior of Group 20 players relative to No Star players. The variable "Accumulated Purchases" is the number of times a player made a purchase between levels $\ell = 21$ and $\ell = 39$. The variable "Indicator of Purchase" is instead a dummy equal to one if "Accumulated Purchases" is positive. The sample includes all 7,409 players in Group 20 and 39,102 in No Star who reached level 20 with a positive star gap, case in which players in Group 20 faced a pay-gate but those in No Star did not. Columns "No Star" and "Group 20" report the means and standard errors, respectively, for the players in No Star and Group 20. Column "No Star−Group 20" reports the difference in mean and associated t-test between the two previous columns. The t-test is calculated assuming unequal variances.

Table A-1: T-test on Non-Pay-Gate Purchases: Group 20 VS No Star

	(1) Group 40		(2) Group 20		(3) Group 40−Group 20	
	mean	SE	mean	SE	diff.	t-test
Accumulated Purchases	0.056	0.521	0.053	0.506	0.003	0.330
Indicator of Purchase	0.021	0.143	0.023	0.149	-0.002	-0.840
Observations	8252		7409		15661	

Notes: This table presents evidence on whether facing an additional pay-gate affects the future non-pay-gate purchasing behavior of Group 20 players relative to No Star players. The variable "Accumulated Purchases" is the number of times a player made a purchase between levels $\ell = 21$ and $\ell = 39$. The variable "Indicator of Purchase" is instead a dummy equal to one if "Accumulated Purchases" is positive. The sample includes all 7,409 players in Group 20 and 8,252 in Group 40 who reached level 20 with a positive star gap, case in which players in Group 20 faced a pay-gate but those in Group 40 did not. Columns "Group 40" and "Group 20" report the means and standard errors, respectively, for the players in Group 40 and Group 20. Column "No Star−Group 20" reports the difference in mean and associated t-test between the two previous columns. The t-test is calculated assuming unequal variances.

Table A-2: T-test on Non-Pay-Gate Purchases: Group 20 VS Group 40

Next, we compare non-pay-gate purchasing behaviors on the basis of a non-parametric Kolmogorov-Smirnov test. The Kolmogorov-Smirnov test detects whether the distribution of a variable differs between two samples. On the basis of the same variables and samples as above, Tables A-3 and A-2 report the Kolmogorov-Smirnov test results. For each variable, the first row assesses whether the players in No Star (or in Group 40) have smaller values than the players in Group 20. The second row instead performs the opposite comparison, assessing whether the players in Group 20 have smaller values than the players in No Star (or in Group 40). The third row "Combined K-S" is the overall test, which is the maximum between the previous two rows. In all cases and for both variables, we observe very high p-values suggesting the absence of any significant difference between the players in No Star (or in Group 40) and those in Group 20. Overall, based on our t-tests and Kolmogorov-Smirnov tests, we do not find evidence that facing the additional pay-gate at $t = 20$ affects players' non-pay-gate purchasing behavior between levels $\ell = 21$ and $\ell = 39$.

	Largest diff.	p-value
Accumulated Purchases		
No Star	0	1
Group 20	0,001121	0,984466
Combined K-S	0,001121	1
Indicator of Purchases		
No Star	0	1
Group 20	0,000667	0,994474
Combined K-S	0,000667	1

Notes: This table presents evidence on whether facing an additional pay-gate affects the future non-pay-gate purchasing behavior of Group 20 players relative to No Star players. The variable "Accumulated Purchases" is the number of times a player made a purchase between levels $\ell = 21$ and $\ell = 39$. The variable "Indicator of Purchase" is instead a dummy equal to one if "Accumulated Purchases" is positive. The sample includes all 7,409 players in Group 20 and 39,102 in No Star who reached level 20 with a positive star gap, case in which players in Group 20 faced a pay-gate but those in No Star did not. Columns "Largest diff." and "p-value" report, respectively, the largest difference and associated p-value for each row on the basis of the Smirnov-Kolmogorov test (where zero means "no difference"). The first row considers the largest difference between the players in No Star and in Group 20. The second row considers the largest difference between the players in Group 20 and in No Star. The last row considers the largest overall difference.

Table A-3: Kolmogorov-Smirnov test on Non-Pay-Gate Purchases: Group 20 VS No Star

	Largest Diff.	p-value
Accumulated Purchases		
Group 40	0,001967	0,970253
Group 20	-0,00125	0,987925
Combined K-S	0,001967	1
Indicator of Purchases		
Group 40	0,001967	0,970253
Group 20	0	1
Combined K-S	0,001967	1

Notes: This table presents evidence on whether facing an additional pay-gate affects the future non-pay-gate purchasing behavior of Group 20 players relative to Group 40 players. The variable "Accumulated Purchases" is the number of times a player made a purchase between levels $\ell = 21$ and $\ell = 39$. The variable "Indicator of Purchase" is instead a dummy equal to one if "Accumulated Purchases" is positive. The sample includes all 7,409 players in Group 20 and 8,252 in Group 40 who reached level 20 with a positive star gap, case in which players in Group 20 faced a pay-gate but those in Group 40 did not. Columns "Largest diff." and "p-value" report, respectively, the largest difference and associated p-value for each row on the basis of the Smirnov-Kolmogorov test (where zero means "no difference"). The first row considers the largest difference between the players in Group 40 and in Group 20. The second row considers the largest difference between the players in Group 20 and in Group 40. The last row considers the largest overall difference.

Table A-4: Kolmogorov-Smirnov test on Non-Pay-Gate Purchases: Group 20 VS Group 40

As a final piece of evidence, we estimate the correlation between pay-gate purchases and non-pay-gate purchases for players in Group 40 (i.e., those who play the standard version of the game). In Table A-5, we regress measures of non-pay-gate purchases on a dummy variable equal to one if the player unlocked pay-gate t by purchasing a key. In the first column, we use as dependent variable the number of non-pay-gate purchases made by the player in the nineteen levels between the two subsequent pay-gates t and $t+20$. In the second column, we instead use a dummy equal to one if the first dependent variable is greater than zero (i.e., the player makes at least one non-pay-gate purchase in these nineteen levels between pay-gate t and $t+20$). Both regressions include player fixed effects and dropout-level fixed effects (these control for the specific levels at which players are observed to drop out of the game). Because of the player fixed effects, the sample includes only the players observed to reach at least two pay-gates with a positive star gap. In both regressions, the estimated coefficient is slightly positive but largely non-significant. Both regressions suggest that once we control for a player's propensity to make in-app purchases (i.e., player fixed effects), pay-gate t purchases do not affect subsequent non-pay-gate purchases between levels $\ell = t+1$ and $\ell = t+19$.

	Nb. purchases in $t < \ell < t + 19$	Any purchase in $t < \ell < t + 19$
Purchase key at t	0.123	0.0256
	(0.178)	(0.0252)
Player FE	Yes	Yes
Dropout-level FE	Yes	Yes
Observations	2,438	2,438
Nb. Players	1,170	1,170

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table presents evidence regarding the relationship between pay-gate purchases (i.e., purchases of keys to unlock pay-gates) and non-pay-gate purchases (i.e., other in-app purchases) for the players in Group 40 (i.e., those who play the standard version of the game). Each column presents OLS estimates of a measure of non-pay-gate purchases on a dummy variable equal to one if the player unlocked pay-gate t by purchasing a key. The first column reports estimation results for the first measure of non-pay-gate purchases: the number of non-pay-gate purchases made by the player in the nineteen levels between the two subsequent pay-gates t and $t + 20$. The second column instead reports estimation results for the second measure: a dummy equal to one if the first measure is greater than zero (i.e., the player makes at least one non-pay-gate purchase in these nineteen levels between pay-gate t and and $t + 20$). Both regressions include player fixed effects and dropout-level fixed effects (these control for the specific levels at which players drop out of the game). Because of the player fixed effects, the sample includes the players in Group 40 observed to reach at least two pay-gates with a positive star gap. Standard errors are clustered at the player level.

Table A-5: Correlation Between Non-Pay-Gate Purchases and Pay-Gate Purchases

B Appendix: kNN Estimator of Model (4)

B.1 Theory

This description of the kNN procedure we use to estimate each of the two binary choice models in (4) is based on Altman (1992).

We observe an i.i.d. sample of data $\{X_i, Y_i\}$ for $i = 1, ..., N$, where X_i is a vector of explanatory variables and Y_i is a binary dependent variable taking values in $\{0, 1\}$. The objective is to estimate the probabilities $Pr(Y_i = 1|X_i)$ and $Pr(Y_i = 0|X_i) = 1 - Pr(Y_i = 1|X_i)$ associated with explanatory variables X_i without making parametric assumptions. To this end, we select a neighborhood $N(X_i)$ of points, with cardinality $k = |N(X_i)|$ around X_i and estimate the sample counter-parts of these probabilities as

$$\widehat{Pr}(Y_i = 1|X_i) = \frac{\sum_{k \in N(X_i)} Y_k}{k}. \tag{11}$$

The neighborhood of points $N(X_i)$ for each observed value of X_i depends on two features. First, the size of the neighborhood denoted generically by the integer $k \in [1, N]$. Second, the distance between any two points X_j and X_s, denoted by $d_{j,s}$, is calculated using a metric. Examples of such metrics include the Euclidean distance ($d_{j,s} = (X_j - X_s)(X_j - X_s)'$), the Mahalanobis distance ($d_{j,s} = (X_j - X_s)V^{-1}(X_j - X_s)'$ where V is the covariance matrix of the matrix X which stacks the vectors of differences), or more generally the Minkowski distance ($d_{j,s} = \left[\sum_{s=1}^{N} |X_j - X_s|^p\right]^{1/p}$ for $p \in \mathbb{N}$).[23] To make variables comparable, we standardize (by subtracting the mean and dividing by the standard deviation) each explanatory variable in X_i. This makes our analysis robust to scale and location distortions.

B.2 Implementation

To select a sufficient number of neighbors k and an appropriate metric d, we search through various possible combinations. In particular, we follow the approach taken in Mitchell (1997) and discussed in Mullin and Sukthankar (2000) by selecting the combination that provides the smallest 5-fold cross validation loss based on the Mean Squared Error $MSE(d, k) = \sum_{j=1}^{N} \left(Y_i - \hat{Y}_j\right)$, where $\hat{Y}_j$ is the model's predicted outcome. The final distance (d), number of neighbors (k), and 10-fold cross validation error rate for equation (4) are reported in Tables A-6 and A-7. We observe a 10-fold cross validation error rate between 15% and 30% which suggests the models are predicting relatively well the underlying probabilities.

[23] In our implementation, other metrics include the correlation distance, the hamming distance, the cosine distance, the Chebychev, the Jaccard distance, and the Spearman distance.

	Observations	k	Distance	Error Rate
Pay-gate 40	7,812	62	Standardized Euclidean	14%
Pay-gate 60	1,433	55	Standardized Euclidean	21%
Pay-gate 80	129	10	Hamming	23%

Table A-6: kNN Estimation : $Pr_{i,t}(sg_{i,t} > 0|i \rightarrow t)$ for Group 40

	Observations	k	Distance	Error Rate
Pay-gate 40	43,660	256	Minkowski	18%
Pay-gate 60	5,205	1,996	Cityblock	28%
Pay-gate 80	917	374	Correlation	14%

Table A-7: kNN Estimation : $Pr_{i,t}(i \rightarrow t + 20)$ for Group 40

B.3 Validation

To validate our kNN estimates, we first compare their accuracy with respect to the underlying data. In Figure A-6, we show that the kNN estimates match, on average, the empirical transition probabilities of equation (4). Moreover, as a sanity check, we display in Figure A-7 the relationship across pay-gates between the estimated probabilities at pay-gate 40 and at pay-gate 60. As expected, players who are more likely to reach pay-gate 40 are also more likely to reach pay-gate 60 (conditional on unlocking pay-gate 40) and those who are more likely to have a positive star gap at pay-gate 40 (conditional on reaching pay-gate 40) are also more likely to have a positive star gap at pay-gate 60 (conditional on reaching pay-gate 60).

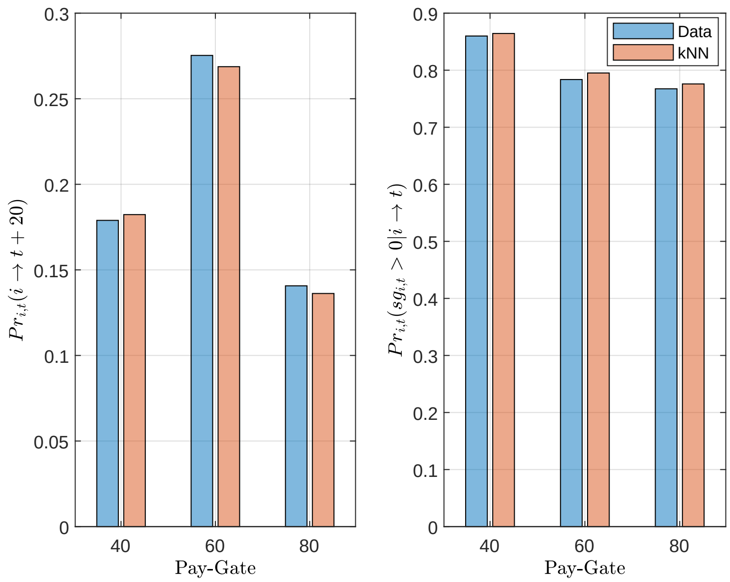

Notes: This figure displays the average probability of reaching the next pay-gate (left panel) and of having a positive star gap conditional on having reached a pay-gate (right panel). These probabilities are displayed based on the data (in blue) and based on the kNN estimates (in red) (described in the main text around equation (4) and above in this Appendix). The sample used is made of all players in Group 40 who have cleared the previous pay-gates.

Figure A-6: Comparing Observed and kNN Estimates of Probabilities in Model (4)

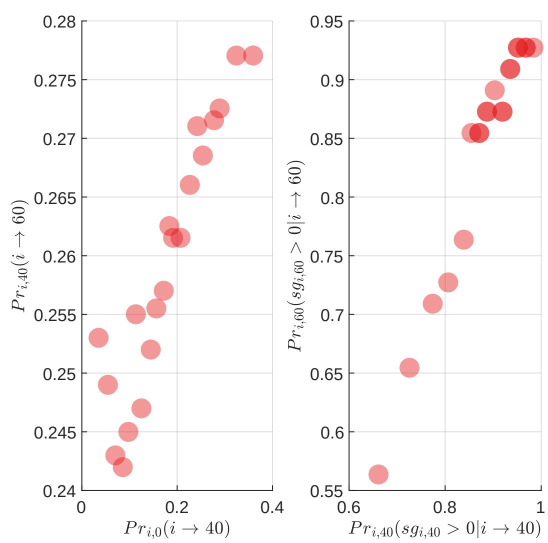

Notes: This figure displays the relationship between the probabilities at pay-gate 40 and 60 of reaching the next pay-gate (left panel) and having a positive star gap (right panel). These binned scatter plots rely on the kNN estimates described in the main text around equation (4) and above in this Appendix. The sample includes all players in Group 40.

Figure A-7: Binned scatter plot of kNN Estimates of model (4) across Pay-gates

C Appendix: Demand Estimates

C.1 First Step Estimates, Equation (8)

In this Appendix, we report the first step estimates of equation (8) and then assess the robustness of our instrument by considering two alternatives.

Figure A-8 shows that the distribution of the instrument $Z_{i,t}$ based on equation (7) has fat tails, confirming the presence of wide sample variation.

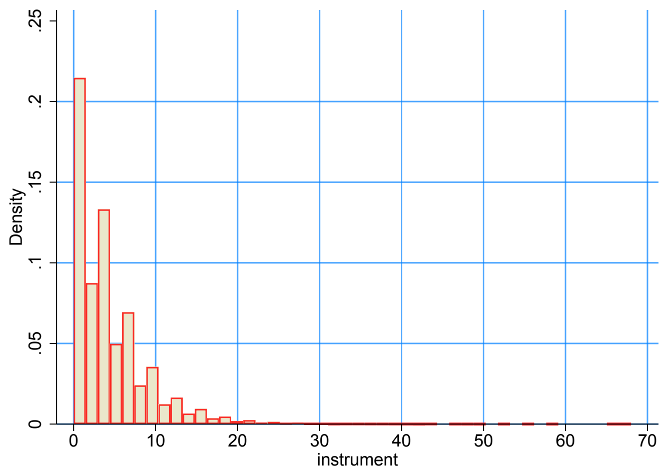

Notes: This histogram displays the dispersion of the instrument $Z_{i,t}$ (based on equation (7)) used to estimate equation (8). The sample includes all observations in which a player from Group "No Star" faced a pay-gate.

Figure A-8: Histogram of the Instrument ($Z_{i,t}$)

Figure A-9 displays a kink in the probability of a non pay-gate purchase when the player nears the necessary score cut-off, clarifying the type of variation leveraged by the proposed instrument.

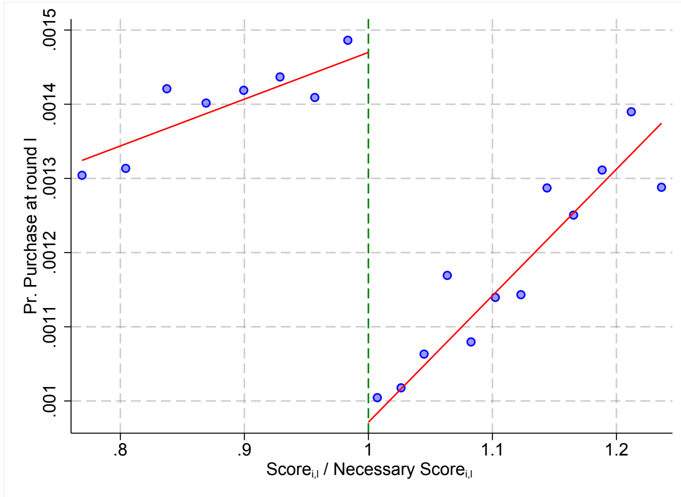

Notes: This binned scatter plot displays the relationship between the probability of a purchase at a level ℓ against the instrument $Z_{i,t}$. The instrument $Z_{i,t}$ is constructed on the basis of equation (7) but adapted for visual inspection as the ratio between the player's score and the minimum required to pass the level attempted. On the y-axis, we plot the probability of making a purchase (outside of a pay-gate) for different values of the instrument on the x-axis. The sample includes all observations of all groups excluding observations corresponding to pay-gates.

Figure A-9: Binned Scatter Plot of Purchase Probability on the Instrument

Figure A-10 confirms the intuition of the instrument: the left panel shows that the more often a player marginally failed a level, the more likely she is to face a lower effective price at the following pay-gate. The right panel illustrates that this is the result of players being pushed to purchase virtual coins to obtain additional lives or moves while trying to clear those challenging levels they marginally failed.

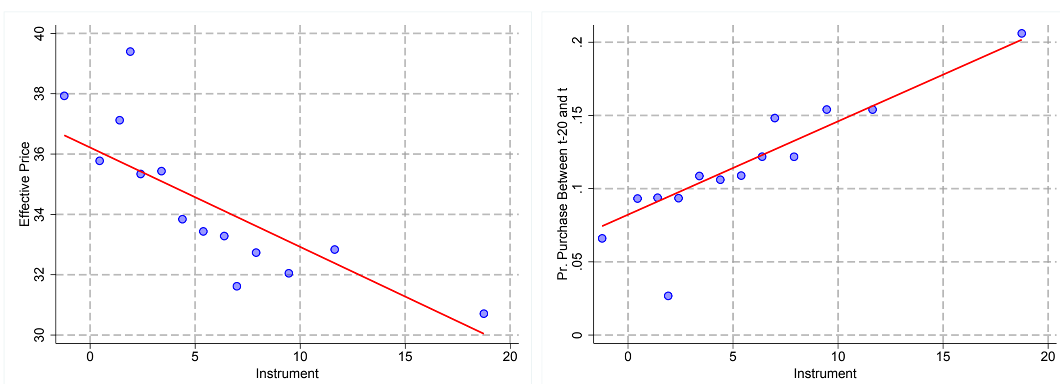

Notes: These binned scatter plots display (left panel) the relationship between the average effective price $p_{i,t}$ against the instrument $Z_{i,t}$ and (right panel) the probability of purchasing virtual coins before reaching the next pay-gate against the instrument $Z_{i,t}$. The instrument $Z_{i,t}$ is constructed on the basis of equation (7). On the y-axis, we plot residualized averages (i.e, the average residual from a regression on pay-gate fixed effects) for various values of the instrument on the x-axis. The sample includes all observations in which players of Group "No Star" faced a pay-gate.

Figure A-10: Binned Scatter Plot of Effective Price on the Instrument

Table A-8 reports the first step estimates of equation (8), which confirms the intuitive patterns from Figure A-10. Conditional on controlling for pay-gate fixed effects and player-specific characteristics X_i (including i's ability), there is a negative and highly statistically significant relationship between the number of times a player marginally failed a round of the game ($Z_{i,t}$) and the effective price ($p_{i,t}$). In particular, each marginal failure is found to lower the effective price by 0.457 virtual coins. The strength of the instrument is reflected in the large F-statistics. All in all, this Table suggests the instrument to be highly informative.

	(1) Effective Price	(2) Effective Price
Instrument $Z_{i,t}$ (δ)	-0.329***	-0.457***
	(0.0355)	(0.0359)
Pay-gate fixed effects (ζ_t)	Yes	Yes
Player-specific characteristics X_i (γ)	No	Yes
Observations	44,385	44,385
Num. of players	37,025	37,025
F-statistic	941.9	179.4

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table reports estimates of equation (8), a regression of the instrument $Z_{i,t}$ based on equation (7) on the effective price $p_{i,t}$. In the first column, we control for pay-gate fixed effects, while in the second we also control for player-specific characteristics X_i (described in Section 3.1). The specification in the second column is the one we use to construct the control function (equation (9)) for the estimation of mixed logit (10). The sample includes all observations in which a player in the Group "No Star" faced a pay-gate: 37,025 different players who faced a total of 44,385 pay-gates. Standard errors are clustered at the player level.

Table A-8: First-Step Estimation, Equation (8)

We now repeat the first step estimation of equation (8) by using two alternatives instruments. We denote by $\mathrm{Z}^{(1)}_{i,t}$ our first alternative instrument, which also counts levels the player marginally cleared:

$$\mathrm{Z}^{(1)}_{i,t} = \sum_{\ell=1}^{t} 1\left(0.95 \times \text{Necessary Score}_{\ell} < \text{Score}_{i,\ell} < 1.05 \times \text{Necessary Score}_{\ell}\right) \tag{12}$$

and by $\mathrm{Z}^{(2)}_{i,t}$ the second alternative instrument, which decreases the threshold below which a player is considered to have marginally failed to pass a given level:

$$\mathrm{Z}^{(2)}_{i,t} = \sum_{\ell=1}^{t} 1\left(0.90 \times \text{Necessary Score}_{\ell} < \text{Score}_{i,\ell} < \text{Necessary Score}_{\ell}\right). \tag{13}$$

Estimation results are reported in Table A-9, which broadly confirm the robustness of the instrument $Z_{i,t}$ to alternative specifications. Varying the definition of the instrument does not qualitatively affect the negative and significant relationship with $p_{i,t}$, as also confirmed by the stability of the F-statistic across regressions. The magnitude of the estimated coefficient halves as a consequence of doubling the length of the interval used for the alternative instruments compared to $Z_{i,t}$.

	(1) Effective Price	(2) Effective Price	(3) Effective Price
Instrument $Z_{i,t}$ (δ)	-0.457***		
	(0.0359)		
Alternative instrument $Z^{(1)}_{i,t}$		-0.267***	
		(0.0196)	
Alternative instrument $Z^{(2)}_{i,t}$			-0.231***
			(0.0189)
Pay-gate fixed effects (ζ_t)	Yes	Yes	Yes
Player-specific characteristics X_i (γ)	Yes	Yes	Yes
Observations	44,385	44,385	44,385
Num. of players	37,025	37,025	37,025
F-statistic	179.4	181.3	178.8

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table reports additional estimation results of equation (8), using alternative definitions of the instrument as in equations (12) and (13). In the first column, we report estimates using the basic definition of the instrument $Z_{i,t}$, while in the second and third we report estimates using the alternative definitions of the instrument $Z^{(1)}_{i,t}$ and $Z^{(2)}_{i,t}$. The sample includes all observations in which a player in the Group "No Star" faced a pay-gate: 37,025 different players who faced a total of 44,385 pay-gates. Standard errors are clustered at the player level.

Table A-9: Robustness of First Step Estimation, Equation (8)

C.2 Second Step Estimates, Equation (10)

Here we report the second step estimates of mixed logit model (10) and then assess their robustness using two alternative instruments described by equations (12) and (13) in Table A-11.

Table A-10 presents our main estimates of the parameters in equation (10). In terms of the utility of purchasing a key using virtual coins ($\text{buy}_{i,t} = 1$), we observe that the constant price coefficient α is negative and statistically significant. The coefficient θ on $\mu_{i,t}$, where $\theta\mu_{i,t}$ is the control function and $\mu_{i,t}$ is estimated using (8) (see Appendix C.1), is positive and significant ($t \approx 11.28$)—confirming the presence of endogeneity in the effective prices.

We asses the robustness of these estimates using two alternative instruments. These instruments are described by equations (12) and (13) in Appendix C.1. For each of these instruments, we re-estimate both the first step equation (8) and the mixed logit model (10). Table A-11 reports our estimates. Compared with the estimates presented in Table A-10, we do not observe any change in the signs across the different specifications for coefficients that are statistically significant. Both the constant associated with the price coefficient (α) and the control function (θ) are within a 95% confidence interval based on the estimates of Table A-10 (i.e, respectively $[-1.01; -0.68]$ and $[-0.51; -0.36]$) suggesting limited differences across specifications. We conclude that our model estimates are robust to alternative specifications of the instrument for effective prices.

Table A-10: Second Stage Estimates

Variable	Coefficient	Standard Error
Purchase ($\text{buy}_{i,t} = 1$)		
Intercepts (δ_1)		
Constant (δ_1)	64,620	6,365
Intercepts Pay-Gate Shifters ($\delta_{1,t}$)		
Pay-Gate 60 ($\delta_{1,60}$)	-4,680	0,628
Pay-Gate 80 ($\delta_{1,80}$)	4,054	0,132
Demographics (β_1)		
Australia and New Zealand	-0,452	0,043
Southern Asia	-1,455	0,170
Southern Europe	-0,112	0,088
Sub Saharan Africa	-3,830	0,370
West Asia	-3,315	0,372
Eastern Asia	0,258	0,014
Eastern Europe	-0,302	0,022
Latin American and Caribbean	2,266	0,189
Northern Africa	-5,614	0,536
Northern America	-2,489	0,256
Northern Europe	-1,976	0,184
Other	-4,925	0,568
South Eastern Asia	0,334	0,052
Control Function (θ)	0,440	0,039
Log(GDP per Capita)	-1,768	0,179
iOS7	-2,225	0,126
iPad	1,487	0,116
Jailbroken	0,490	0,062
Ability	-3,929	0,350
Effective Price Intercept (α)		
Constant (α)	-0,852	0,083
Effective Price Pay-Gate Shifters		
Pay-Gate 60 (α_{60})	0,036	0,007
Pay-Gate 80 (α_{80})	0,053	0,008
Effective Price Demographic (π)		
Australia and New Zealand	0,017	0,001
Southern Asia	0,019	0,002
Southern Europe	0,004	0,001

Table A-10: Second Stage Estimates

Variable	Coefficient	Standard Error
Sub Saharan Africa	-0,014	0,000
West Asia	-0,015	0,001
Eastern Asia	-0,007	0,000
Eastern Europe	0,020	0,003
Latin American and Caribbean	0,022	0,002
Northern Africa	-0,222	0,017
Northern America	0,007	0,001
Northern Europe	0,005	0,000
Other	0,005	0,000
South Eastern Asia	0,018	0,002
Log(GDP per Capita)	0,020	0,002
iOS7	-0,002	0,000
iPad	-0,001	0,000
Jailbroken	-0,018	0,002
Ability	0,006	0,001
Ask a Friend ($\text{buy}_{i,t} = 2$)		
Intercepts (δ_2)		
Constant (δ_2)	-43,650	6,696
Intercepts Pay-Gate Shifters ($\delta_{2,t}$)		
Pay-Gate 60 ($\delta_{2,60}$)	-8,078	1,100
Pay-Gate 80 ($\delta_{2,80}$)	-4,662	0,612
Demographics (β_2)		
Australia and New Zealand	-4,307	0,611
Southern Asia	-4,406	0,379
Southern Europe	-1,531	0,206
Sub Saharan Africa	-7,681	0,724
West Asia	-3,578	0,456
Eastern Asia	-6,415	0,865
Eastern Europe	-6,631	0,958
Latin American and Caribbean	-1,158	0,108
Northern Africa	0,998	0,312
Northern America	-5,635	0,810
Northern Europe	0,334	0,055
Other	-3,531	0,281
South Eastern Asia	3,459	0,544

Table A-10: Second Stage Estimates

Variable	Coefficient	Standard Error
Log(GDP per Capita)	4,458	0,668
iOS7	1,850	0,248
iPad	-1,509	0,223
Jailbroken	1,447	0,109
Ability	-0,348	0,036
Covariance Matrix		
$\sigma_{\delta_{2,i}}$	10,025	1,491
$\rho_{\delta_{2,i},\delta_{1,i}}$	0,432	0,081
$\sigma_{\delta_{1,i}}$	3,241	0,485
$\rho_{\delta_{2,i},\alpha_i}$	-0,043	0,004
$\rho_{\delta_{1,i},\alpha_i}$	-0,035	0,003
σ_{α_i}	-0,001	0,000

Notes: This table reports estimation results of equation (10). Standard errors are calculated using the method by Karaca-Mandic and Train (2003) to account for the two-step nature of the estimation procedure. The estimation procedure is described in Section 5.2 and uses 100 Halton draws per player as detailed by Bhat (2003). Each variable is defined in Section 3.1. The reference region is Western Europe. The sample includes all observations in which a player in the Group "No Star" faced a pay-gate: 37,025 different players who faced a total of 44,385 pay-gates.

Table A-11: Robustness of Second Stage Estimates

	Alternative Instrument $Z^{(1)}_{i,t}$		Alternative Instrument $Z^{(2)}_{i,t}$	
Variable	Coefficient	Standard Error	Coefficient	Standard Error
Purchase ($\text{buy}_{i,t} = 1$)				
Intercepts (δ_1)				
Constant (δ_1)	69,907	4,767	57,527	4,329
Intercepts Pay-Gate Shifters ($\delta_{1,t}$)				
Pay-Gate 60 ($\delta_{1,60}$)	-5,653	0,470	-4,481	0,175
Pay-Gate 80 ($\delta_{1,80}$)	4,268	0,645	4,384	1,058
Demographics (β_1)				
Australia and New Zealand	-0,677	0,055	-0,643	0,119
Southern Asia	-1,748	0,109	-1,424	0,094
Southern Europe	-0,419	0,098	-0,181	0,059
Sub Saharan Africa	-4,257	0,252	-3,464	0,244
West Asia	-3,613	0,196	-2,716	0,128
Eastern Asia	0,061	0,035	-0,014	0,047
Eastern Europe	-0,312	0,064	-0,319	0,054
Latin American and Caribbean	2,141	0,571	2,024	0,337
Northern Africa	-5,905	0,542	-4,822	0,268
Northern America	-2,645	0,030	-1,983	0,095
Northern Europe	-2,168	0,106	-1,809	0,101
Other	-5,296	0,645	-4,658	0,241
South Eastern Asia	0,277	0,192	0,223	0,055
Control Function (θ)	0,483	0,060	0,418	0,047
Log(GDP per Capita)	-1,746	0,281	-1,406	0,055
iOS7	-2,452	0,239	-2,209	0,344
iPad	1,536	0,197	1,321	0,079
Jailbroken	0,601	0,070	0,513	0,037
Ability	-4,324	0,614	-3,644	0,351
Effective Price Intercept (α)				
Constant (α)	-0,913	0,006	-0,742	0,042
Effective Price Pay-Gate Shifters				
Pay-Gate 60 (α_{60})	0,045	0,001	0,030	0,005
Pay-Gate 80 (α_{80})	0,058	0,001	0,038	0,007
Effective Price Demographic (π)				
Australia and New Zealand	0,021	0,001	0,018	0,001
Southern Asia	0,019	0,011	0,011	0,001
Southern Europe	0,004	0,004	0,002	0,002

Table A-11: Robustness of Second Stage Estimates

	Alternative Instrument $Z^{(1)}_{i,t}$		Alternative Instrument $Z^{(2)}_{i,t}$	
Variable	Coefficient	Standard Error	Coefficient	Standard Error
Sub Saharan Africa	-0,012	0,018	-0,012	0,003
West Asia	-0,015	0,005	-0,013	0,002
Eastern Asia	-0,006	0,010	-0,005	0,001
Eastern Europe	0,025	0,000	0,019	0,000
Latin American and Caribbean	0,025	0,005	0,017	0,001
Northern Africa	-0,297	2599	-0,623	5660246
Northern America	0,009	0,002	0,009	0,000
Northern Europe	0,006	0,001	0,005	0,001
Other	0,010	0,005	0,003	0,004
South Eastern Asia	0,017	0,003	0,011	0,001
Log(GDP per Capita)	0,019	0,010	0,015	0,000
iOS7	-0,003	0,000	-0,003	0,000
iPad	0,000	0,001	0,000	0,001
Jailbroken	-0,017	0,002	-0,015	0,001
Ability	0,007	0,002	0,004	0,000
Ask a Friend ($\text{buy}_{i,t} = 2$)				
Intercepts (δ_2)				
Constant (δ_2)	-65,547	28,772	-51,045	7,597
Intercepts Pay-Gate Shifters ($\delta_{2,t}$)				
Pay-Gate 60 ($\delta_{2,60}$)	-10,966	2,887	-8,690	0,982
Pay-Gate 80 ($\delta_{2,80}$)	-5,752	0,919	-4,531	0,219
Demographics (β_2)				
Australia and New Zealand	-6,424	2,581	-5,020	0,879
Southern Asia	-5,733	2,715	-4,081	3,227
Southern Europe	-2,080	0,645	-1,700	0,268
Sub Saharan Africa	-10,525	5,170	-9,185	5,016
West Asia	-5,090	1,872	-3,992	0,572
Eastern Asia	-8,789	2,244	-7,007	0,943
Eastern Europe	-9,015	1,214	-7,260	0,814
Latin American and Caribbean	-1,515	0,203	-1,301	0,335
Northern Africa	1,655	0,567	0,678	1,019
Northern America	-8,222	3,084	-6,457	1,013
Northern Europe	0,433	0,095	0,315	0,032
Other	-4,723	1,013	-3,657	0,426
South Eastern Asia	5,072	1,795	4,104	0,617

Table A-11: Robustness of Second Stage Estimates

	Alternative Instrument $Z_{i,t}^{(1)}$		Alternative Instrument $Z_{i,t}^{(2)}$	
Variable	Coefficient	Standard Error	Coefficient	Standard Error
Log(GDP per Capita)	6,554	2,743	5,109	0,724
iOS7	2,622	0,876	2,114	0,343
iPad	-2,234	0,858	-1,752	0,318
Jailbroken	1,459	0,672	1,501	0,682
Ability	-0,469	0,176	-0,377	0,078
Covariance Matrix				
$\sigma_{\delta_{2,i}}$	15,174	5,826	11,948	2,230
$\rho_{\delta_{2,i},\delta_{1,i}}$	1,756	0,103	1,639	0,084
$\sigma_{\delta_{1,i}}$	-3,346	1,741	1,719	0,883
$\rho_{\delta_{2,i},\alpha_i}$	-0,061	0,004	-0,053	0,013
$\rho_{\delta_{1,i},\alpha_i}$	0,029	0,026	-0,013	0,024
σ_{α_i}	0,001	0,003	0,000	0,007

Notes: This table reports estimation results of equation (10). Standard errors are calculated using the method by Karaca-Mandic and Train (2003) to account for the two-step nature of the estimation procedure. The estimation procedure is described in Section 5.2 and uses 100 Halton draws per player as detailed by Bhat (2003). Each variable is defined in Section 3.1. The reference region is Western Europe. The sample includes all observations in which a player in the Group "No Star" faced a pay-gate: 37,025 different players who faced a total of 44,385 pay-gates.

D Appendix: Price Elasticities and Counterfactual Simulations

D.1 Formulae

In this Appendix, we detail the formulae used to compute all our model predictions and simulations.

Price Elasticity of Demand. For player i at pay-gate t and given η_i, we refer to the multinomial logit formula as:

$$\text{MNL}_{i,t}(\eta_i) = \frac{\exp(V_{1,i,t}(\eta_i))}{1 + \exp(V_{1,i,t}(\eta_i)) + \exp(V_{2,i,t}(\eta_i))},$$

where, as described in Section 5.2, $V_{1,i,t}(\eta_i) = \delta_1 + \delta_{1,t} + \delta_{1,i} + X_i\beta_1 - (\alpha + \alpha_t + \alpha_i + X_i\pi)p_{i,t} + \theta\mu_{i,t}$ includes control function $\theta\mu_{i,t}$, based on (7), (8), and (9), to account for the potential endogeneity of $p_{i,t}$, and $V_{2,i,t}(\eta_i) = \delta_2 + \delta_{2,t} + \delta_{2,i} + X_i\beta_2$. Then, mixed logit model (10) implies the following price elasticity of demand:

$$\begin{aligned} &\frac{p_{i,t}}{Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t})} \frac{\partial Pr_{i,t}(\text{buy}_{i,t} = 1|\text{lock}_{i,t} = 1, p_{i,t})}{\partial p_{i,t}} = \\ &\frac{-p_{i,t}(\alpha + \alpha_t + X_i\pi)}{\int \text{MNL}_{i,t}(\eta_i)\phi(\eta_i|\Sigma)d\eta_i} \int \alpha_i \text{MNL}_{i,t}(\eta_i)\left(1 - \text{MNL}_{i,t}(\eta_i)\right)\phi(\eta_i|\Sigma)d\eta_i, \end{aligned} \tag{14}$$

where $\phi(\cdot|\Sigma)$ is the normal density of η_i in (6) with Σ denoting its variance-covariance matrix.

Per-Player Expected Revenue. Here we derive the formulae we use to compute the per-player expected revenue at level 0 in all our simulations with the exception of Figure 13 (which we instead discuss in the next sub-section). We calculate the per-player expected revenue at level 0 from pay-gate 40 (and none of the next pay-gates) as:

$$\mathbb{E}_{i,0}[R_{i,40}(p_{i,40})] = Pr_{i,0}(i \to 40, sg_{i,40} > 0) \times Pr_{i,40}(\text{buy}_{i,40} = 1|\text{lock}_{i,40} = 1, p_{i,40}) \times p_{i,40}$$

where $Pr_{i,0}(i \to 40, sg_{i,40} > 0)$ is i's probability of reaching pay-gate 40 with a positive star gap given that the player is at the beginning of the game, at level 0. This can be simply expressed in terms of the estimated probabilities in equation (4) as $Pr_{i,0}(i \to 40, sg_{i,40} > 0) = Pr_{i,40}(i \to 40, sg_{i,40} > 0)$, given that $t = 40$ is the first pay-gate i can encounter in the game. Per-player expected revenue at level 0 from pay-gate 60 (and none of the next pay-gates) is equal to:

$$\mathbb{E}_{i,0}[R_{i,60}(p_{i,40}, p_{i,60})] = Pr_{i,0}(i \to 60, sg_{i,60} > 0) \times Pr_{i,60}(\text{buy}_{i,60} = 1|\text{lock}_{i,60} = 1, p_{i,60}) \times p_{i,60}.$$

where $Pr_{i,0}(i \to 60, sg_{i,60} > 0)$ is i's probability of reaching pay-gate 60 with a positive star gap given that she is at level 0. We can again express this in terms of the estimated probabilities in models (4) and (10) as:

$$Pr_{i,0}(i \to 60, sg_{i,60} > 0) = Pr_{i,60}(i \to 60, sg_{i,60} > 0) \times Pr_{i,40}(i \to 40) \times \\ \left[Pr_{i,40}(sg_{i,40} > 0 | i \to 40) \times (1 - Pr_{i,40}(\text{buy}_{i,40} = 0 | \text{lock}_{i,40} = 1, p_{i,40}) + (1 - Pr_{i,40}(sg_{i,40} > 0 | i \to 40))\right],$$

where $Pr_{i,60}(i \to 60, sg_{i,60} > 0)$ denotes i's probability of reaching pay-gate $t = 60$ with a positive star gap given that she unlocked pay-gate $t = 40$, and so on for the other probabilities. Similarly, we can write the per-player expected revenue at level 0 from pay-gate 80 as:

$$\mathbb{E}_{i,0}[R_{i,80}(p_{i,40}, p_{i,60}, p_{i,80})] = Pr_{i,0}(i \to 80, sg_{i,80} > 0) \times Pr_{i,80}(\text{buy}_{i,80} = 1 | \text{lock}_{i,80} = 1, p_{i,80}) \times p_{i,80}.$$

where $Pr_{i,0}(i \to 80, sg_{i,80} > 0)$ is i's probability of reaching pay-gate 80 with a positive star gap given that she is at level 0. This can be expressed in terms of the estimated probabilities in models (4) and (10) as:

$$Pr_{i,0}(i \to 80, sg_{i,80} > 0) = Pr_{i,80}(i \to 80, sg_{i,80} > 0) \times \\ Pr_{i,40}(i \to 40) \times \left[Pr_{i,40}(sg_{i,40} > 0 | i \to 40) \times (1 - Pr_{i,40}(\text{buy}_{i,40} = 0 | \text{lock}_{i,40} = 1, p_{i,40}) + (1 - Pr_{i,40}(sg_{i,40} > 0 | i \to 40))\right] \times \\ Pr_{i,60}(i \to 60) \times \left[Pr_{i,60}(sg_{i,60} > 0 | i \to 60) \times (1 - Pr_{i,60}(\text{buy}_{i,60} = 0 | \text{lock}_{i,60} = 1, p_{i,60}) + (1 - Pr_{i,60}(sg_{i,60} > 0 | i \to 60))\right].$$

Finally, we calculate the per-player expected revenue from player i at level 0 (from all pay-gates) given effective prices $p_i = (p_{i,40}, p_{i,60}, p_{i,80})$ as:

$$R_{i,0}(p_i) = \mathbb{E}_{i,0}[R_{i,40}(p_i) + R_{i,60}(p_i) + R_{i,80}(p_i)]. \tag{15}$$

This is the central expression at the basis of our counterfactual simulations, i.e. what the firm maximizes when choosing effective prices, and the main focus of the simulation method described in Appendix D.2.

Revenue Decomposition in Figure 13. Here we derive the formulae used in Figure 13. While the computation of $Pr_{i,40}(\text{buy}_{i,40} = 1 | \text{lock}_{i,40} = 1, p_{i,40}) \times p_{i,40}$ is immediate from the estimates of model (10), the computation of $(1 - Pr_{i,40}(\text{buy}_{i,40} = 0 | \text{lock}_{i,40} = 1, p_{i,40})) \times R_{i,60}(p_{i,60} | p_{i,80})$ requires the calculation of $R_{i,60}(p_{i,60} | p_{i,80})$. This in turn can be expressed as a function of the probabilities estimated in models (4) and (10) as:

$$R_{i,60}(p_{i,60} | p_{i,80}) = Pr_{i,40}(i \to 60, sg_{i,60} > 0) \times Pr_{i,60}(\text{buy}_{i,60} = 1 | \text{lock}_{i,60} = 1, p_{i,60}) \times p_{i,60} \\ + Pr_{i,60}(i \to 60) \times \left[Pr_{i,60}(sg_{i,60} > 0 | i \to 60) \times (1 - Pr_{i,60}(\text{buy}_{i,60} = 0 | \text{lock}_{i,60} = 1, p_{i,60}) + (1 - Pr_{i,60}(sg_{i,60} > 0 | i \to 60))\right] \\ \times Pr_{i,80}(i \to 80, sg_{i,80} > 0) Pr_{i,80}(\text{buy}_{i,80} = 1 | \text{lock}_{i,80} = 1, p_{i,80}) \times p_{i,80},$$

Per-Player Consumer Surplus. Here we derive the formulae we use to compute changes in per-player consumer surplus at level 0 in all our simulations. The derivations follow closely those for the per-player expected revenue above and here we rely on some of the objects defined there. We calculate the per-player consumer surplus at level 0 from pay-gate 40 (and none of the next pay-gates) as:

$$\mathbb{E}_{i,0}[CS_{i,40}(p_{i,40})] = Pr_{i,0}(i \to 40, sg_{i,40} > 0) \times CS_{i,40}(\text{lock}_{i,40} = 1|X_i, p_{i,40}).$$

Based on standard formulae for mixed logit models (Train, 2009), the per-player consumer surplus at level 40 from pay-gate 40 is given by:

$$CS_{i,40}(\text{lock}_{i,40} = 1|X_i, p_{i,40}) = \mathscr{C}_{i,40} + \mathbb{E}_{i,40}\left(\frac{\ln(1+\exp(V_{1,i,40}(\eta_i)) + \exp(V_{2,i,40}(\eta_i)))}{(\alpha + \alpha_{40} + \alpha_i + X_i\pi)}\right)$$

where $\mathscr{C}_{i,40}$ is an unknown player-specific constant. Similarly, per-player consumer surplus at level 0 from pay-gate 60 (and none of the next pay-gates) is equal to:

$$\mathbb{E}_{i,0}[CS_{i,60}(p_{i,40}, p_{i,60})] = Pr_{i,0}(i \to 60, sg_{i,60} > 0) \times CS_{i,60}(\text{lock}_{i,60} = 1|X_i, p_{i,60}),$$

where

$$CS_{i,60}(\text{lock}_{i,60} = 1|X_i, p_{i,60}) = \mathscr{C}_{i,60} + \mathbb{E}_{i,60}\left(\frac{\ln(1+\exp(V_{1,i,60}(\eta_i)) + \exp(V_{2,i,60}(\eta_i)))}{(\alpha + \alpha_{60} + \alpha_i + X_i\pi)}\right).$$

In turn, we can write the per-player consumer surplus at level 0 from pay-gate 80 as:

$$\mathbb{E}_{i,0}[CS_{i,80}(p_{i,40}, p_{i,60}, p_{i,80})] = Pr_{i,0}(i \to 80, sg_{i,80} > 0) \times CS_{i,80}(\text{lock}_{i,80} = 1|X_i, p_{i,80}).$$

where

$$CS_{i,80}(\text{lock}_{i,80} = 1|X_i, p_{i,80}) = \mathscr{C}_{i,80} + \mathbb{E}_{i,80}\left(\frac{\ln(1+\exp(V_{1,i,80}(\eta_i)) + \exp(V_{2,i,80}(\eta_i)))}{(\alpha + \alpha_{80} + \alpha_i + X_i\pi)}\right).$$

We calculate the per-player consumer surplus at level 0 (from all pay-gates) given effective prices $p_i = (p_{i,40}, p_{i,60}, p_{i,80})$ as:

$$CS_{i,0}(p_i) = \mathbb{E}_{i,0}[CS_{i,40}(p_{i,40}) + CS_{i,60}(p_{i,60}) + CS_{i,80}(p_{i,80})]. \tag{16}$$

Finally, for any given two vectors of effective prices p_i and p_i', we compute the associated change in per-player consumer surplus simply as:

$$\Delta CS_{i,0}(p_i, p_i') = CS_{i,0}(p_i) - CS_{i,0}(p_i').$$

Note that this difference at the player-level has the important advantage of removing, for each i, the unknown constants $\mathscr{C}_{i,40}$, $\mathscr{C}_{i,60}$, and $\mathscr{C}_{i,80}$.

D.2 Simulation Method

Here we describe our simulation procedure both in the case of the observed pricing strategy chosen by the firm and in the case of the counterfactual pricing strategies we investigate. In general, to simulate the model, one needs to specify each player's effective price $p_{i,t}$ and corresponding residual $\mu_{i,t}$ from equation (8) at each possible pay-gate t.

Observed Pricing Strategy. Even in the case of the observed pricing strategy, because some players may have dropped out before reaching pay-gate 80, we cannot back out $(p_{i,t}, \mu_{i,t})$ directly from the data for all players and pay-gates. To address this missing data problem, we follow the approach by Jacobi and Sovinsky (2016) and treat the unobserved $(p_{i,t}, \mu_{i,t})$ as random effects to be integrated over their empirical distribution.

We are interested in simulating the firm's expected revenue (16) from player i before they start to play (at $t = 0$) for any given vector of effective prices $p_i = (p_{i,40}, p_{i,60}, p_{i,80})$ and corresponding residuals from equation (8), $\mu_i = (\mu_{i,40}, \mu_{i,60}, \mu_{i,80})$. To stress the dependence on both p_i and μ_i, due to the control function in mixed logit model (10), we extend the notation of per-player expected revenue (16) to explicitly account also for μ_i, $R_{i,0}(p_i, \mu_i)$. For those players who did not reach pay-gate t, we cannot directly back out $(p_{i,t}, \mu_{i,t})$ but assume that it follows the same empirical distribution $\hat{F}_{p,\mu,t}$ as among those players observed to reach pay-gate t. In particular, we compute the joint distribution $\hat{F}_{p,\mu,t}$ as $\hat{F}_{p,t}\hat{F}_{\mu|p,t}$, the product of the unconditional distribution of $p_{i,t}$ and the conditional distribution of $\mu_{i,t}$ given $p_{i,t}$. We then approximate the per-player expected revenue

$$\mathbb{E}_{\hat{F}_{p,\mu}}\left[R_{i,0}(p_i, \mu_i)\right] = \int R_{i,0}(p_i, \mu_i) d\hat{F}_{p,\mu,40}(p_{i,40}, \mu_{i,40}) d\hat{F}_{p,\mu,60}(p_{i,60}, \mu_{i,60}) d\hat{F}_{p,\mu,80}(p_{i,80}, \mu_{i,80}) \tag{17}$$

by taking 10,000 draws of (p_d, μ_d) from $\hat{F}_{p,\mu} = \hat{F}_{p,\mu,40}\hat{F}_{p,\mu,60}\hat{F}_{p,\mu,80}$ for each i and computing the average:

$$\hat{\mathbb{E}}_{\hat{F}_{p,\mu}}\left[R_{i,0}(p_i, \mu_i)\right] = \frac{1}{10,000} \times \sum_{d=1}^{10,000} R_{i,0}(p_d, \mu_d), \tag{18}$$

where $R_{i,0}(p_i, \mu_i)$ is derived above in Appendix D.1, equation (16).

Counterfactual Pricing Strategy. In each of the counterfactual pricing strategies described in Section 6.2 and detailed below, the firm chooses effective prices so to maximize the sum of per-player expected revenue $R_{i,0}(p_i, \mu_i)$ in equation (16) (see Appendix D.1) across players subject to some constraints. To simulate these counterfactuals, we assume that μ_i is also unobserved to firm and that its distribution is invariant to the specific pricing strategy used. More precisely, we assume that $\mu_{i,t}$ follows the same empirical distribution $\hat{F}_{\mu,t}$ as among those players observed to reach pay-gate t and that the firm uses $\hat{F}_{\mu,t}$ to form expectations with respect to $\mu_{i,t}$. We then approximate the per-player expected revenue for a given vector of effective prices p_i

$$\mathbb{E}_{\hat{F}_\mu}\left[R_{i,0}(p_i,\mu_i)\right] = \int R_{i,0}(p_i,\mu_i)\, d\hat{F}_{\mu,40}(\mu_{i,40})\, d\hat{F}_{\mu,60}(\mu_{i,60})\, d\hat{F}_{\mu,80}(\mu_{i,80}) \tag{19}$$

by taking 10,000 draws of μ_d from $\hat{F}_\mu = \hat{F}_{\mu,40}\hat{F}_{\mu,60}\hat{F}_{\mu,80}$ for each i and computing the average:

$$\hat{\mathbb{E}}_{\hat{F}_\mu}\left[R_{i,0}(p_i,\mu_i)\right] = \frac{1}{10,000} \times \sum_{d=1}^{10,000} R_{i,0}(p_i,\mu_d). \tag{20}$$

Counterfactual Pricing Strategies. Here we detail the optimization problem of the firm in the simulation of each of the counterfactual pricing strategies described in Section 6.2. All counterfactuals are computed for the 44,660 players in Group 40 (those who play the standard version of the game) over pay-gates 40, 60, and 80.

- ☐ Uniform (Optimal) Static Pricing:

$$(p^*, p^*, p^*) = \arg\max_{p} \sum_{i=1}^{44,660} \hat{\mathbb{E}}_{\hat{F}_\mu}\left[R_{i,0}(p,p,p,\mu_i)\right].$$

- ☐ Uniform (Optimal) Dynamic Pricing:

$$(p^*_{40}, p^*_{60}, p^*_{80}) = \arg\max_{p_{40},p_{60},p_{80}} \sum_{i=1}^{44,660} \hat{\mathbb{E}}_{\hat{F}_\mu}\left[R_{i,0}(p_{40},p_{60},p_{80},\mu_i)\right].$$

- ☐ GDP per Capita Static Pricing for players in ventile $G = 1, ..., 20$ of GDP per capita:

$$(p^*_G, p^*_G, p^*_G) = \arg\max_{p_G} \sum_{i\in G} \hat{\mathbb{E}}_{\hat{F}_\mu}\left[R_{i,0}(p_G,p_G,p_G,\mu_i)\right],$$

  where each ventile G gathers 5% of players in terms of the observed distribution of GDP per capita.

- ☐ GDP per Capita Dynamic Pricing for players in ventile $G = 1, ..., 20$ of GDP per capita:

$$(p^*_{G,40}, p^*_{G,60}, p^*_{G,80}) = \arg\max_{p_{G,40},p_{G,60},p_{G,80}} \sum_{i\in G} \hat{\mathbb{E}}_{\hat{F}_\mu}\left[R_{i,0}(p_{G,40},p_{G,60},p_{G,80},\mu_i)\right],$$

  where each ventile G gathers 5% of players in terms of the observed distribution of GDP per capita.

- □ Ability Static Pricing for players in ventile $A = 1, ..., 20$ of ability:

$$(p_A^*, p_A^*, p_A^*) = \arg\max_{p_A} \sum_{i \in A} \hat{\mathbb{E}}_{\hat{F}_\mu} \left[R_{i,0}(p_A, p_A, p_A, \mu_i) \right],$$

  where each ventile A gathers 5% of players in terms of the observed distribution of ability.

- □ Ability Dynamic Pricing for players in ventile $A = 1, ..., 20$ of ability:

$$(p_{A,40}^*, p_{A,60}^*, p_{A,80}^*) = \arg\max_{p_{A,40}, p_{A,60}, p_{A,80}} \sum_{i \in A} \hat{\mathbb{E}}_{\hat{F}_\mu} \left[R_{i,0}(p_{A,40}, p_{A,60}, p_{A,80}, \mu_i) \right],$$

  where each ventile A gathers 5% of players in terms of the observed distribution of ability.

- □ Individual Static Pricing for player $i = 1, ..., 43660$:

$$(q_i^*, q_i^*, q_i^*) = \arg\max_{q_i} \hat{\mathbb{E}}_{\hat{F}_\mu} \left[R_{i,0}(q_i, q_i, q_i, \mu_i) \right].$$

- □ Individual Dynamic Pricing for player $i = 1, ..., 43660$:

$$p_i^* = \arg\max_{p_i} \hat{\mathbb{E}}_{\hat{F}_\mu} \left[R_{i,0}(p_i, \mu_i) \right].$$

For simplicity, we solve each of these optimization problems using a simple grid search over effective prices. For each effective price the firm can choose, we specify a grid with intervals of 5 virtual coins going from 0 to 100, $[0, 5, 10, ..., 95, 100]$. For example, in Uniform (Optimal) Static Pricing this results in 21 possible combinations of effective prices, while in Individual Dynamic Pricing in $21^3 = 9,261$ combinations for each player $i = 1, ..., 43660$. We then evaluate the per-player expected revenue for each combination of effective prices and player and solve the above optimization problems. We do not extend the support of the optimizations above 100 virtual coins as we never found any optimal effective price to be larger than 70 virtual coins (which is also the maximum value we observe in the data). The step size of 5 virtual coins was selected as a trade-off between precision and required computational time.

E Appendix: Model Validation

Our simulation exercises rely on the estimated model's ability to predict player behavior under counterfactual pricing strategies. In this section, we investigate the predictive power of the estimated model in terms of a player's expected revenue. As mentioned at the beginning of Section 6, while we relied on the players in Group 40 for the estimation of model (4) and on those in No Stars for the estimation of model (10), we test the model's predictive power and perform all counterfactual simulations only with respect to the players in Group 40. Players in Group 40 face the default design of the game, which corresponds to our discrete choice model in equation (4).

As can be seen in equation (3), to compute the per-player expected revenue at pay-gate t, we need to know the effective prices they would face at t and later pay-gates. However, for those players that drop out of the game before pay-gate t, we do not observe these effective prices. We address this "missing data" problem as in Jacobi and Sovinsky (2016) and treat the effective prices as another dimension of unobserved heterogeneity to be integrated over when calculating expectations. We assume the true t-specific distribution of $p_{i,t}$ across players, $F_t(p_{i,t})$, can be consistently estimated as the empirical distribution of the observed effective prices at t, $\hat{F}_t(p_{i,t})$, and then integrate per-player expected revenue over $\hat{F}_t(p_{i,t})$ for each i. In Appendix D.1 we report the formulae used to compute per-player expected revenue (used also in the counterfactual simulations) and in Appendix D.2 we describe the details of this simulation procedure (and of the procedure used for the counterfactual simulations). Below we report our validation results comparing the per-player average observed revenue with its counterpart as predicted by the estimated model.

The estimated model is very good at predicting the average observed per-player revenue of \$0.011 (from purchases of keys to unlock pay-gates), delivering a t-test as small as -0.09249. Note that this result is not mechanical, in that the estimated parameters are not chosen to minimize the distance between observed and predicted revenue, but rather the probabilities of models (4) and (10). Importantly for the investigation of price discrimination, Figures A-11 and A-12 illustrate the accuracy of the estimated model in predicting revenues for specific profiles of the observed player-specific characteristics X_i.

Figure A-11 compares the total revenue by geographical region (in \$) as observed in the data against that as predicted by the estimated model across different geographical regions. We calculate observed total revenue as the sum over 43,660 players in Group 40 of the revenue collected at pay-gates 40, 60, and 80 in each geographical region during the 15 days of our sample in 2013. The figure confirms the presence of geographical heterogeneity in total revenue and that the estimated model is good at capturing it. In general, the geographical regions for which more revenue is observed (and so for which we have more observations) are also those for which the estimated model delivers more accurate predictions.

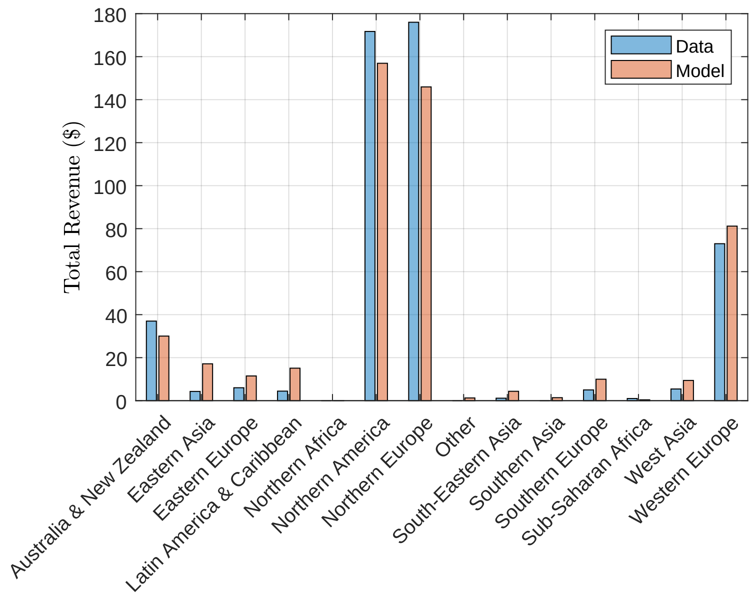

Notes: This figure compares the total revenue by geographical region (in $) as observed in the data against that as predicted by the estimated model across different geographical regions. We calculate the observed total revenue as the sum over 43,660 players in Group 40 of the revenue collected at pay-gates 40, 60, and 80 in each geographical region during the 15 days of our sample in 2013. The expected total revenue as predicted by our estimated model is based on the simulation procedure detailed in Appendices D.1 and D.2. The geographical regions are described in Section 3.1 and Appendix G.

Figure A-11: Observed Against Simulated Total Revenue by Region

Figure A-12 compares the total revenue (in $) as observed in the data against that as predicted by the estimated model across players with different ability (left panel) and from countries with different GDP per capita (right panel). In particular, the left panel reports results by deciles (D1 being the lowest and D10 the highest decile) of ability while the right panel by deciles of log(GDP per capita). We calculate observed total revenue as the sum over 43,660 players in Group 40 of the revenue collected at pay-gates 40, 60, and 80 in each decile during the 15 days of our sample in 2013. The left panel shows that, when it comes to ability, the estimated model does a good job at predicting total revenue for 50% of players, those with ability between the third and the seventh decile. However, it tends to under-predict for most able players (the top three deciles) and to over-predict for the least able ones (the bottom two deciles). The right panel of Figure A-12 confirms that, with the exception of the bottom two deciles of poorest countries, the estimated model is overall good at predicting total revenue in terms of players' log(GDP per capita). As we discuss at the end of Section 6 in the main text, in robustness checks reported in Appendix F.1, we account for these predictive biases by limiting our counterfactual simulations to the sub-sample of players in Group 40 for whom the estimated model has better predictive power (deciles D3-D7 of ability and D3-D10 of GDP per capita).

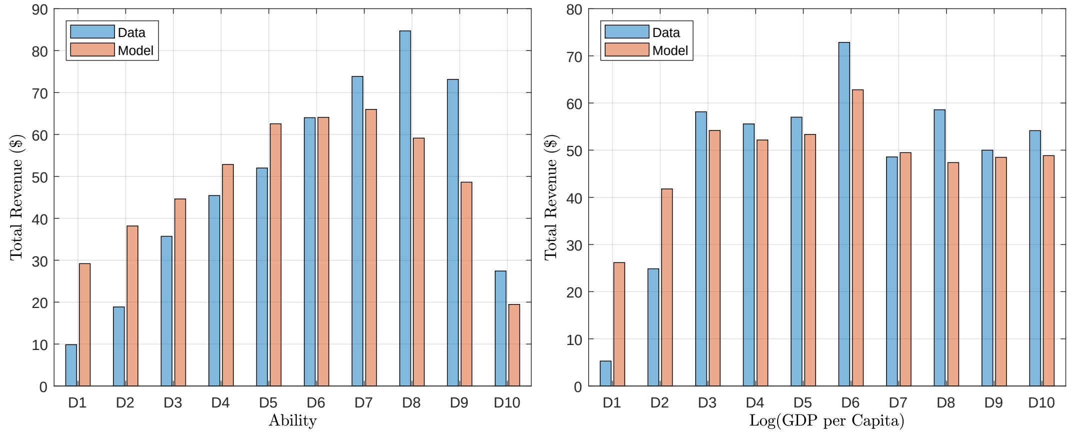

Notes: These figures compare the total revenue (in $) as observed in the data against that as predicted by the estimated model across players with different ability (left panel) and from countries with different GDP per capita (right panel). In particular, the left panel reports results by deciles (D1 is the lowest decile and D10 the highest) of players' ability while the right panel by deciles of log(GDP per capita). Ability and log(GDP per capita) are described in Section 3.1. We calculate the observed total revenue as the sum over 43,660 players in Group 40 of the revenue collected at pay-gates 40, 60, and 80 in each decile during the 15 days of our sample in 2013. The expected total revenue as predicted by our estimated model is based on the simulation procedure detailed in Appendices D.1 and D.2.

Figure A-12: Observed Against Predicted Revenue by Ability and GDP per Capita

F Appendix: Additional Simulation Results

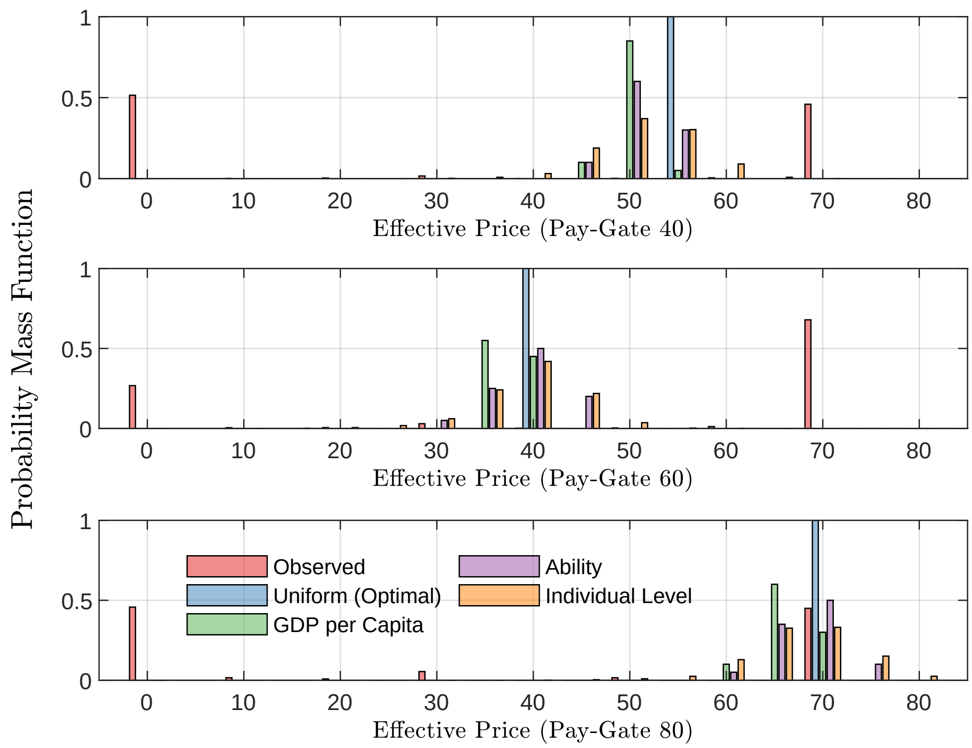

Notes: This figure shows the simulated distribution of effective prices (in virtual coins, where \$1 ≈ 70 virtual coins) across players within each pay-gate for the "dynamic" pricing strategies considered in the right panel of Table 7. All pricing strategies are decreibed in detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices are allowed to change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-13: Distribution of Effective Prices in Dynamic Pricing Strategies

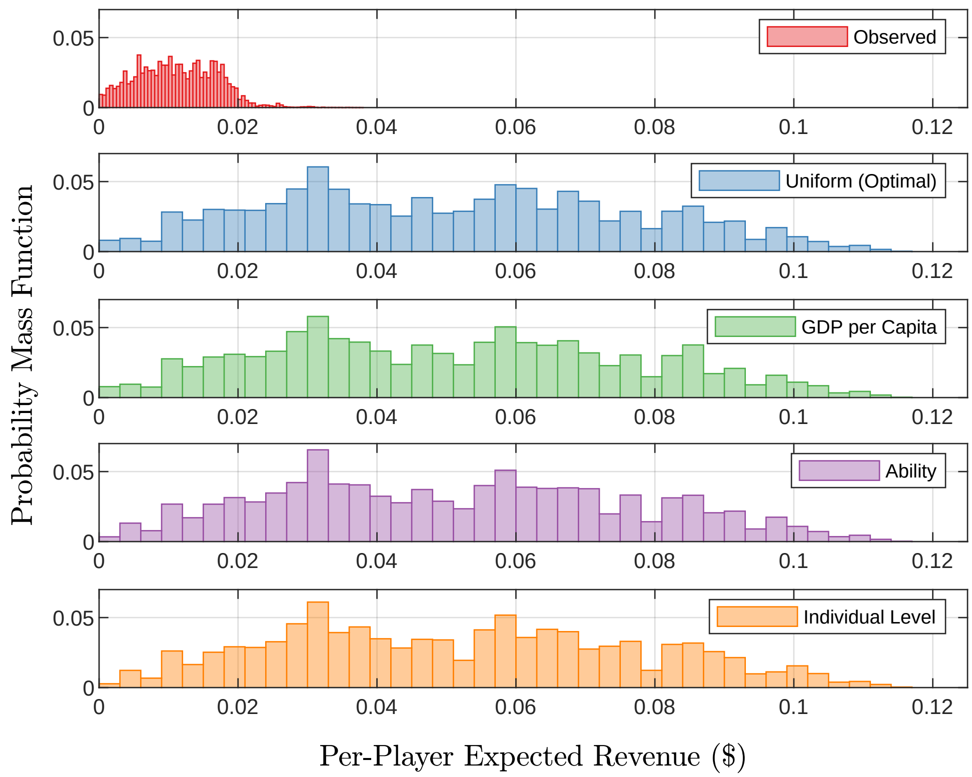

Notes: This figure shows the simulated distribution of per-player expected revenue (in $) across players for the "dynamic" pricing strategies considered in the right panel of Table 7. All pricing strategies are described in detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices are allowed to change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-14: Distribution of Per-Player Expected Revenue, Dynamic Pricing Strategies

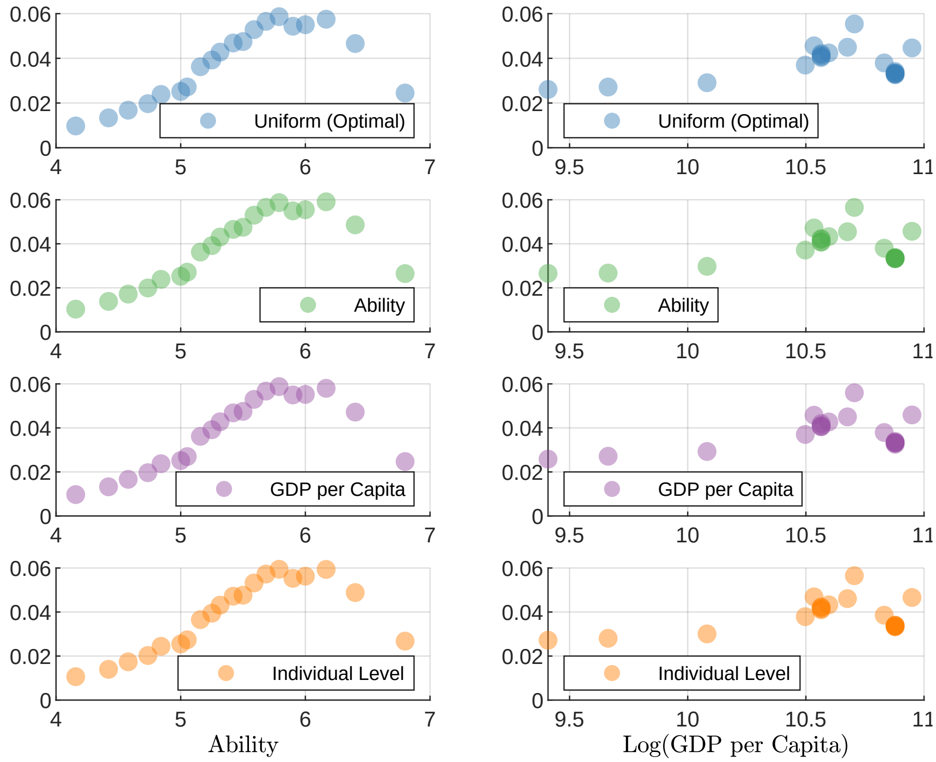

Notes: These binscatters display the average gains in per-player expected revenue (in $) of engaging in the "static" pricing strategies considered in the left panel of Table 7 as opposed to the observed pricing. For each static pricing strategy, we construct 20 groups of players based on players' ability (left panel) and log(GDP per capita) (right panel) and plot the average group-specific difference in per-player expected revenue on the y-axis. The definitions of ability and log(GDP per capita) are provided in Section 3.1. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-15: Δ Per-Player Expected Revenue, Static versus Observed Pricing

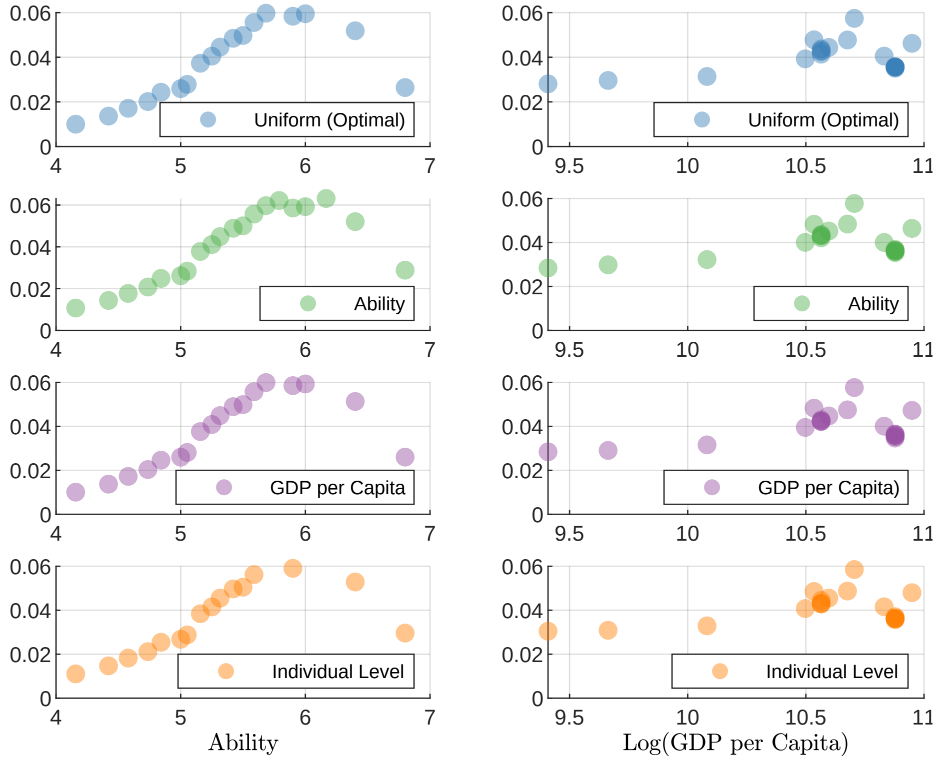

Notes: These binscatters display the average gains in per-player expected revenue (in $) of engaging in the "dynamic" pricing strategies considered in the right panel of Table 7 as opposed to the observed pricing. For each static pricing strategy, we construct 20 groups of players based on players' ability (left panel) and log(GDP per capita) (right panel) and plot the average group-specific difference in per-player expected revenue on the y-axis. The definitions of ability and log(GDP per capita) are provided in Section 3.1. All pricing strategies are described in detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices are allowed to change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-16: Δ Per-Player Expected Revenue, Dynamic versus Observed Pricing

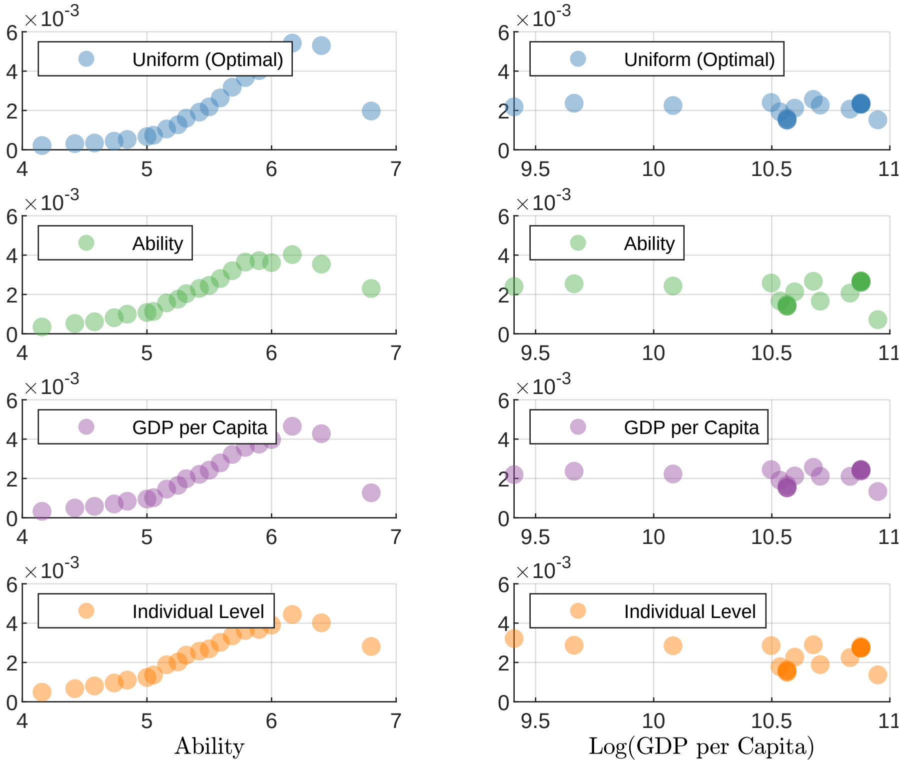

Notes: These binscatters display the average gains in per-player expected revenue (in $) of engaging in the "dynamic" versus the "static" versions of each of the pricing strategies considered in Table 7. For each pricing strategy, we construct 20 groups of players based on players' ability (left panel) and log(GDP per capita) (right panel) and plot the average group-specific difference in per-player expected revenue on the y-axis. The definitions of ability and log(GDP per capita) are provided in Section 3.1. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Static pricing strategies are those in which effective prices do not change among pay-gates. Dynamic pricing strategies are instead those in which effective prices are also allowed to change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-17: Δ Per-Player Expected Revenue, Dynamic versus Static Pricing

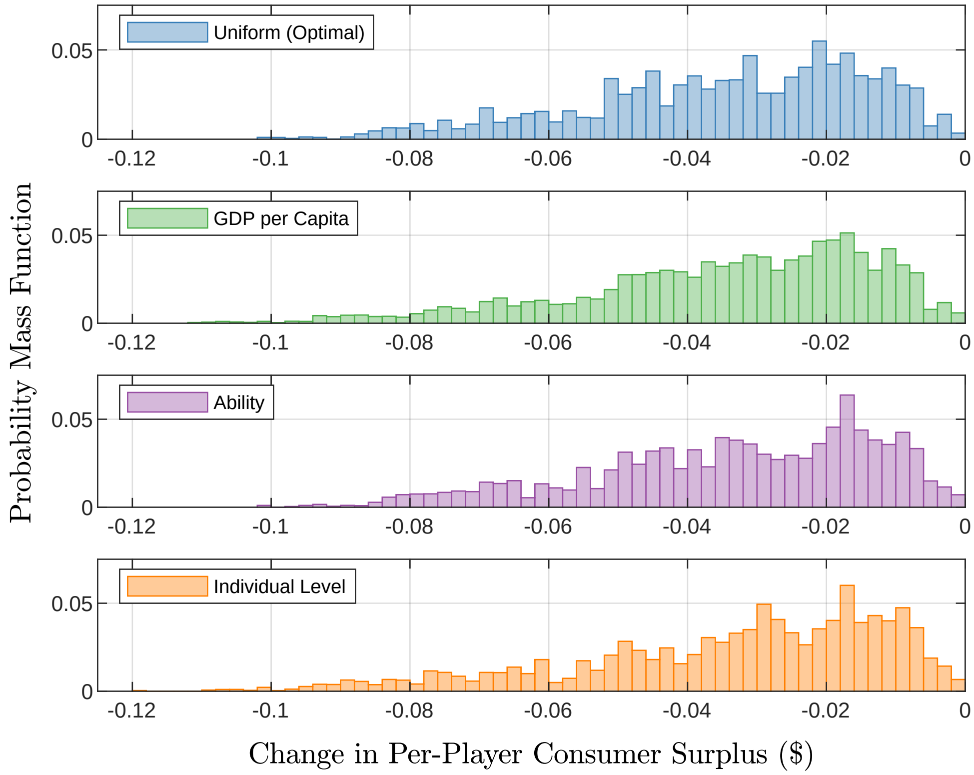

Notes: This figure shows the simulated distribution of changes in per-player consumer surplus (in $) across players for the "dynamic" pricing strategies considered in the right panel of Table 8. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-18: Distribution of Δ Per-Player Consumer Surplus, Dynamic versus Observed Pricing

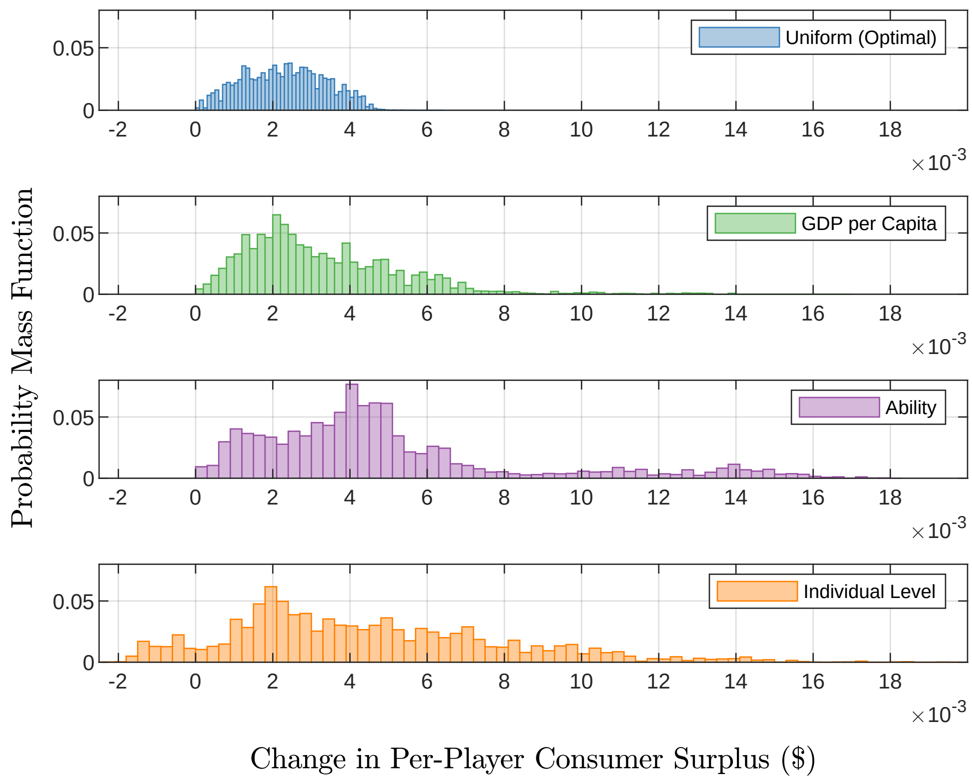

Notes: This figure shows the simulated distribution of changes in per-player consumer surplus (in $) across players of engaging in the "dynamic" versus the "static" versions of each of the pricing strategies considered in Table 8. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices change among pay-gates. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-19: Distribution of Δ Per-Player Consumer Surplus, Dynamic versus Static Pricing

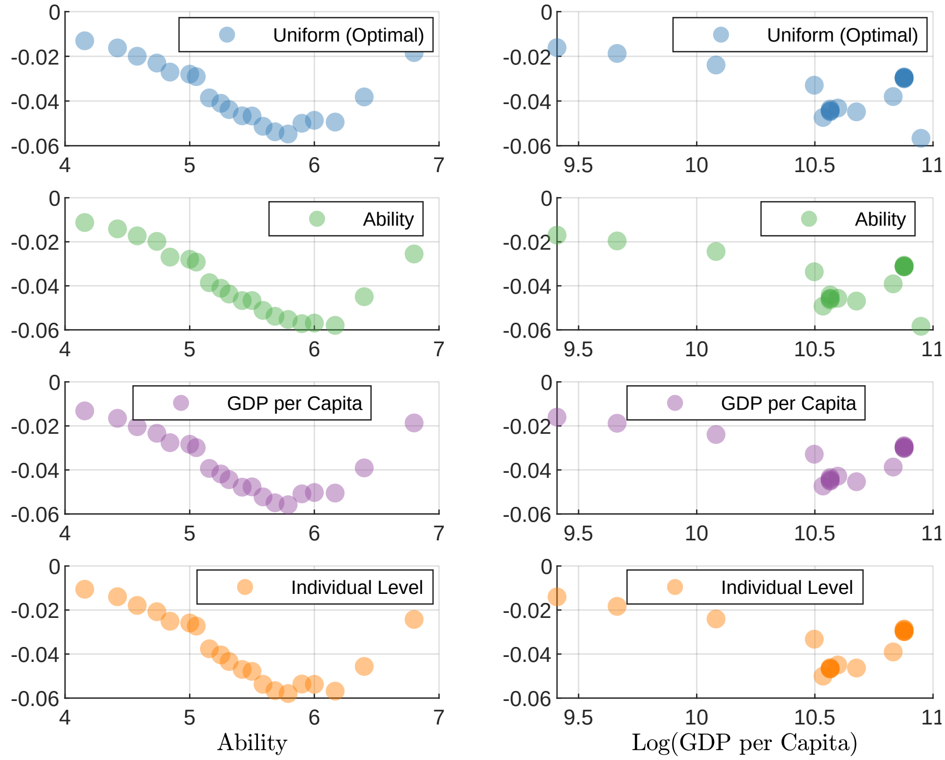

Notes: These binscatters display the average change in per-player consumer surplus (in $) of engaging in the "static" pricing strategies considered in the left panel of Table 8 as opposed to the observed pricing. For each static pricing strategy, we construct 20 groups of players based on players' ability (left panel) and log(GDP per capita) (right panel) and plot the average group-specific difference in per-player expected consumer surplus on the y-axis. The definitions of ability and log(GDP per capita) are provided in Section 3.1. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-20: Δ Per-Player Consumer Surplus, Static versus Observed Pricing

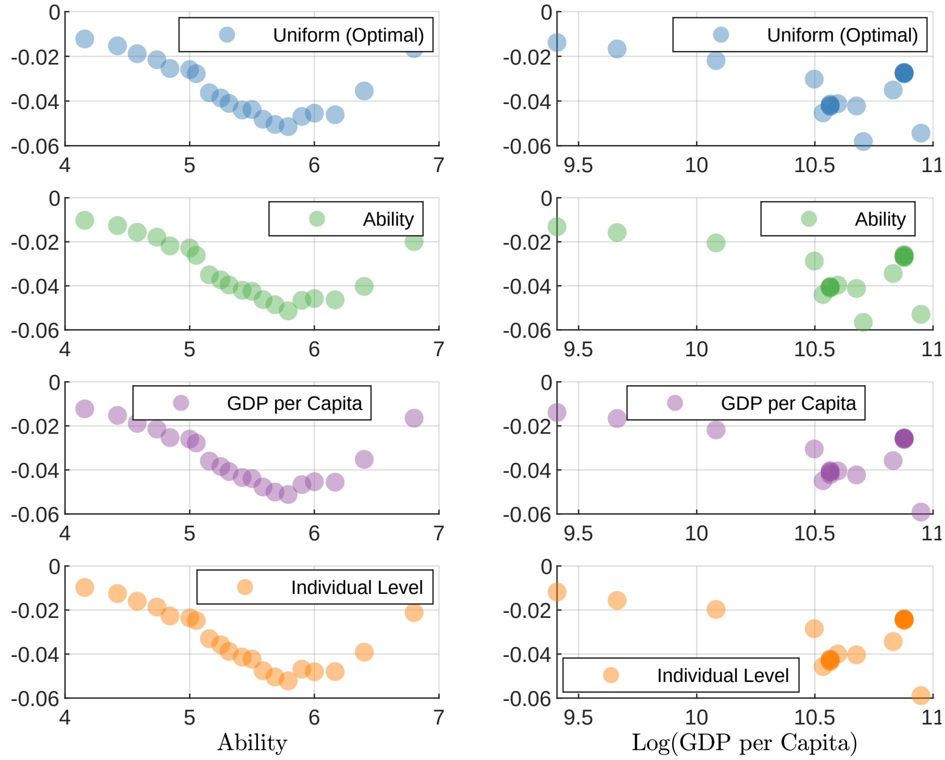

Notes: These binscatters display the average change in per-player consumer surplus (in \$) of engaging in the "dynamic" pricing strategies considered in the right panel of Table 8 as opposed to the observed pricing. For each dynamic pricing strategy, we construct 20 groups of players based on players' ability (left panel) and log(GDP per capita) (right panel) and plot the average group-specific difference in per-player consumer surplus on the y-axis. The definitions of ability and log(GDP per capita) are provided in Section 3.1. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-21: Δ Per-Player Consumer Surplus, Dynamic versus Observed Pricing

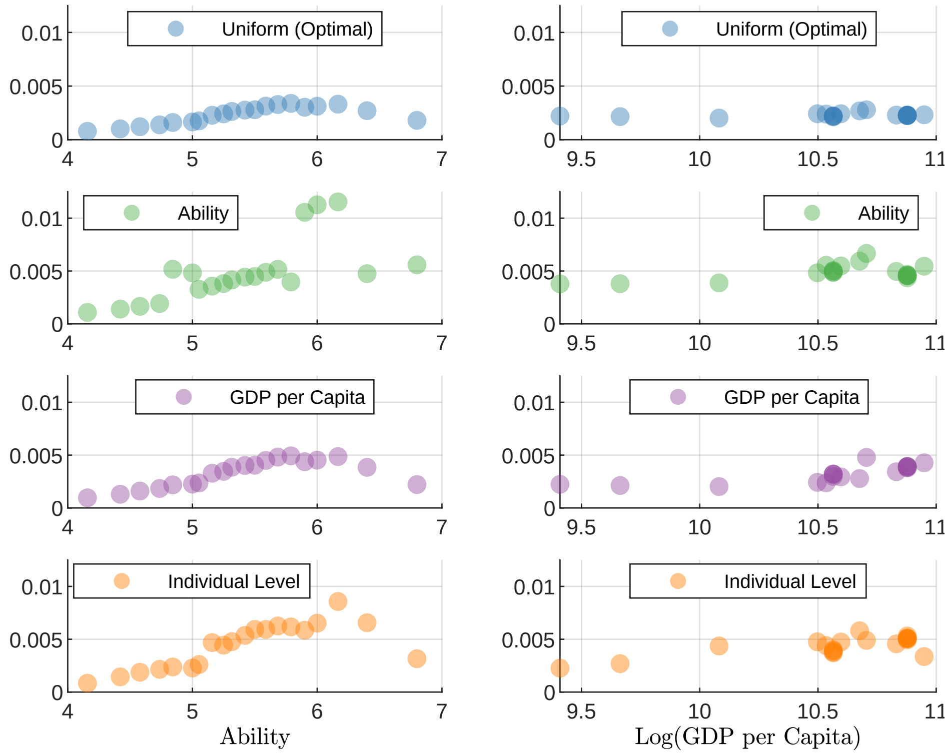

Notes: These binscatters display the average change in per-player consumer surplus (in $) of engaging in the "dynamic" versus the "static" versions of each of the pricing strategies considered in Table 8. For each pricing strategy, we construct 20 groups of players based on players' ability (left panel) and log(GDP per capita) (right panel) and plot the average group-specific difference in per-player consumer surplus on the y-axis. The definitions of ability and log(GDP per capita) are provided in Section 3.1. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices change among pay-gates. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-22: Δ Per-Player Consumer Surplus, Dynamic versus Static Pricing

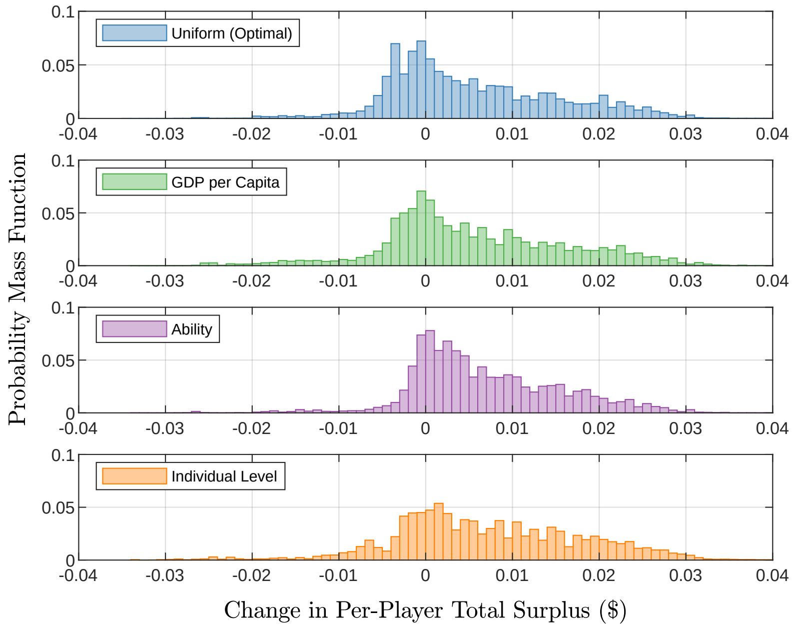

Notes: This figure shows the simulated distribution of changes in per-player total surplus (in $) across players for the "dynamic" pricing strategies considered in the right panel of Table 8 as opposed to the observed pricing. Changes in per-player total surplus are computed as the sum between changes in per-player expected revenues and in per-player consumer surplus. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-23: Distribution of Δ Per-Player Total Surplus, Dynamic versus Observed Pricing

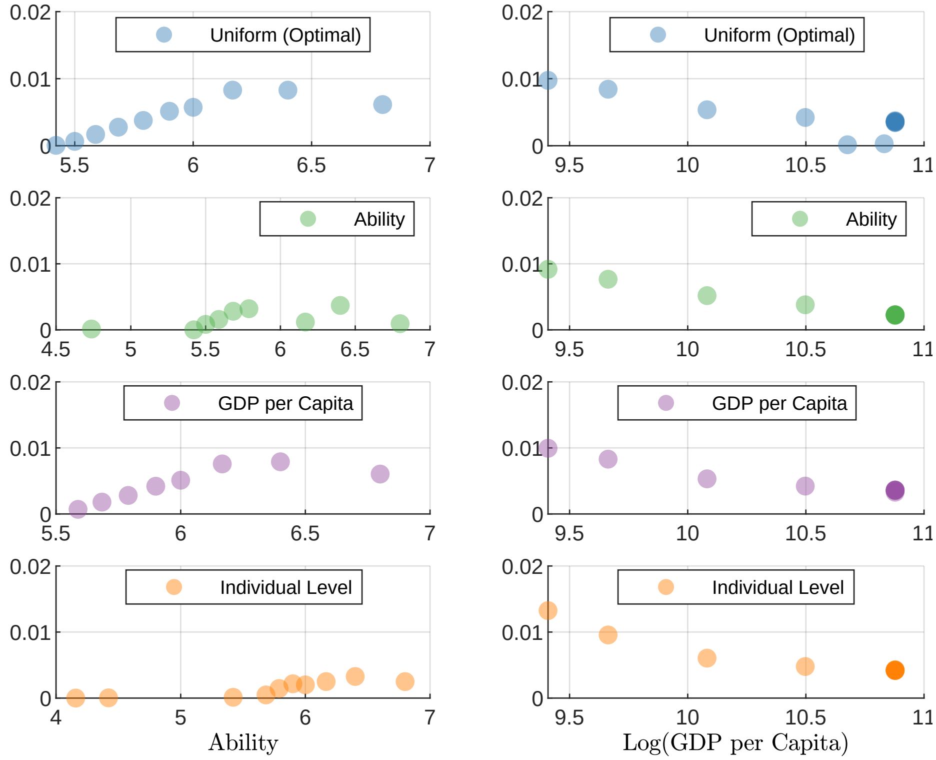

Notes: These binscatters display the average gains in per-player total surplus (in \$) of engaging in the "static" pricing strategies considered in the left panel of Table 8 as opposed to the observed pricing. For each static pricing strategy, we construct 20 groups of players based on players' ability (left panel) and log(GDP per capita) (right panel) and plot the average group-specific difference in per-player expected revenue on the y-axis. The definitions of ability and log(GDP per capita) are provided in Section 3.1. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. Static pricing strategies are those in which effective prices do not change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-24: Δ in Per-Player Total Surplus, Static versus Observed Pricing

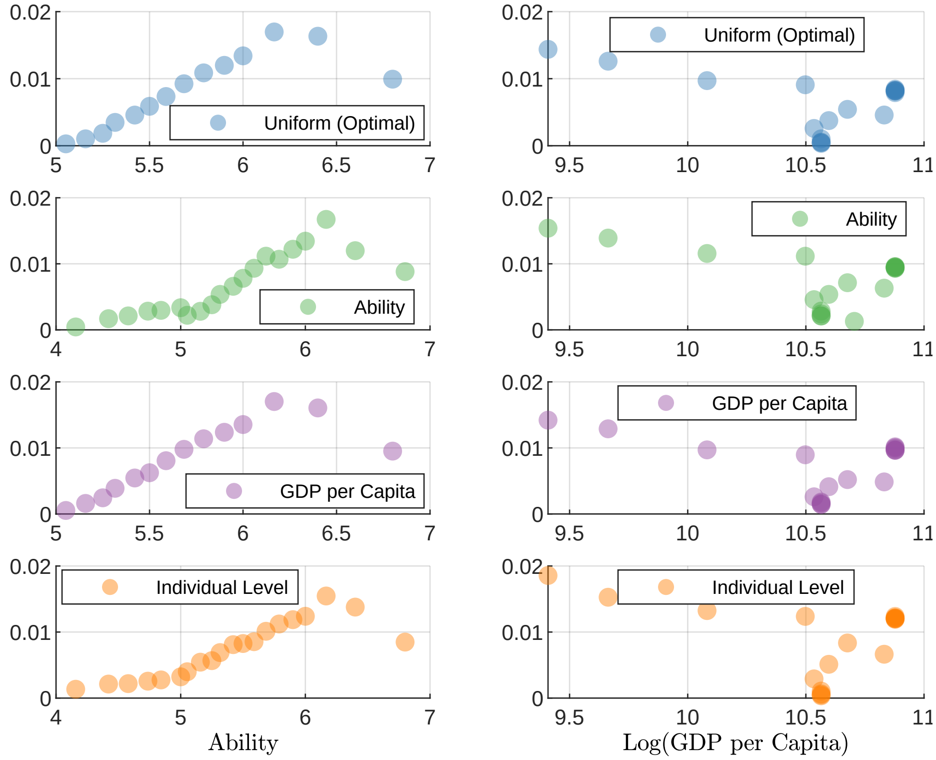

Notes: These binscatters display the average gains in per-player total surplus (in $) of engaging in the "dynamic" pricing strategies considered in the right panel of Table 8 as opposed to the observed pricing. For each dynamic pricing strategy, we construct 20 groups of players based on players' ability (left panel) and log(GDP per capita) (right panel) and plot the average group-specific difference in per-player expected revenue on the y-axis. The definitions of ability and log(GDP per capita) are provided in Section 3.1. All pricing strategies are described in detail in Appendix D.2. Dynamic pricing strategies are those in which effective prices are allowed to change among pay-gates. All simulations are based on our estimates of models (4) and (10) and on the 43,660 players in Group 40 during the 15 days of our sample in 2013. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Figure A-25: Δ in Per-Player Total Surplus, Dynamic versus Observed Pricing

F.1 Robustness Checks

In this Appendix, we repeat all counterfactual simulations limiting the sample of players in Group 40 to those for whom the estimated model displays the best predictive power in terms of expected revenue, namely the players with ability in deciles D3-D7 and GDP per capita in deciles D3-D10. We do this in order to account for the predictive biases of our estimated model as documented in Appendix E. Overall, these robustness checks show no qualitative difference in any of our results and suggest that the predictive biases documented in Appendix E do not play a crucial role in our simulations exercises.

Pricing Strategy	Static Pricing					Dynamic Pricing				
	Effective Price		Per-Player Revenue ($)			Effective Price		Per-Player Revenue ($)		
	mean	s.d.	mean	s.d.	%	mean	s.d.	mean	s.d.	%
Observed	35,566	34,529	0,014	0,120	-	-	-	-	-	-
Uniform (70)	70,000	-	0,028	0,008	98.1%	-	-	-	-	-
Uniform (Optimal)	45,000	-	0,056	0,020	292.4%	10,000	2,450	0,057	0,021	302.3%
GDP per Capita	44,385	1,642	0,056	0,020	292.2%	51,451	11,839	0,058	0,021	304.8%
Ability	45,001	0,075	0,056	0,020	292.4%	53,323	12,059	0,058	0,021	305.1%
Individual Level Pricing	45,603	3,086	0,056	0,020	294.1%	53,418	12,073	0,058	0,021	308.1%

Notes: This table summarizes our counterfactual simulation results in terms of effective prices and per-player expected revenues using a restricted sample. Each row refers to a pricing strategy and summarizes the simulated effective prices chosen by the firm (in virtual coins, where $1 ≈ 70 virtual coins) and the corresponding per-player expected revenues (in $). The columns denoted by "%" report the percentage increase in per-player expected revenue implied by the row pricing strategy with respect to the observed pricing chosen by the firm (i.e., 0% means same average as the observed pricing). All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. The left panel summarizes results for the case in which effective prices do not change among pay-gates (static pricing). The right panel instead summarizes results for the case in which effective prices are allowed to change also among pay-gates (dynamic pricing). All simulations are based on our estimates of models (4) and (10). The sample excludes players in Group 40 who are below the 2nd decile in terms of GDP per Capita, below the third decile in terms of ability, and above the seventh decile in ability. There are 17,719 remaining players. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Table A-12: Robustness Check : Simulation of Effective Prices and Expected Revenue with Restricted Sample

Pricing Strategy	Static Pricing				Dynamic Pricing			
	Δ Consumer Surplus ($)		Δ Total Surplus ($)		Δ Consumer Surplus ($)		Δ Total Surplus ($)	
	mean	s.d.	mean	s.d.	mean	s.d.	mean	s.d.
Uniform (70)	-0,0227	0,0128	-0,0088	0,0115	-	-	-	-
Uniform (Optimal)	-0,0441	0,0187	-0,0025	0,0059	-0,0417	0,0181	0,0014	0,0068
GDP per Capita	-0,0451	0,0202	-0,0035	0,0077	-0,0412	0,0192	0,0022	0,0085
Ability	-0,0441	0,0187	-0,0025	0,0059	-0,0396	0,0178	0,0038	0,0066
Individual Level	-0,0443	0,0215	-0,0025	0,0079	-0,0397	0,0207	0,0042	0,0094

Notes: This table summarizes our counterfactual simulation results in terms of per-player consumer surplus and per-player total surplus, computed as the sum between changes in per-player expected revenue and in per-player consumer surplus using a restricted sample. Each row refers to a pricing strategy and summarizes the simulated change in per-player consumer surplus and in per-player total surplus (both in $) with respect to the observed pricing. All pricing strategies are briefly described in the text and explained in more detail in Appendix D.2. The left panel summarizes results for the case in which effective prices do not change among pay-gates (static pricing). The right panel instead summarizes results for the case in which effective prices are allowed to change also among pay-gates (dynamic pricing). All simulations are based on our estimates of models (4) and (10). The sample excludes players in Group 40 who are below the 2nd decile in terms of GDP per Capita, below the third decile in terms of ability, and above the seventh decile in ability. There are 17,719 remaining players. Details of the formulae and simulation procedures used can be found in Appendices D.1 and D.2.

Table A-13: Robustness Check : Simulation of Consumer Surplus and Total Surplus with Restricted Sample

G Appendix: Data

Table A-14: Countries and Region Assignment

Country	Region
Afghanistan	Southern Asia
Ãland Islands	Northern Europe
Albania	Southern Europe
Algeria	Northern Africa
American Samoa	Other
Andorra	Southern Europe
Angola	Sub-Saharan Africa
Anguilla	Latin America and the Caribbean
Antarctica	Other
Antigua and Barbuda	Latin America and the Caribbean
Argentina	Latin America and the Caribbean
Armenia	Western Asia
Aruba	Latin America and the Caribbean
Australia	Australia and New Zealand
Austria	Western Europe
Azerbaijan	Western Asia
Bahamas	Latin America and the Caribbean
Bahrain	Western Asia
Bangladesh	Southern Asia
Barbados	Latin America and the Caribbean
Belarus	Eastern Europe
Belgium	Western Europe
Belize	Latin America and the Caribbean
Benin	Sub-Saharan Africa
Bermuda	Northern America
Bhutan	Southern Asia
Bolivia (Plurinational State of)	Latin America and the Caribbean
Bonaire, Sint Eustatius and Saba	Latin America and the Caribbean
Bosnia and Herzegovina	Southern Europe
Botswana	Sub-Saharan Africa
Bouvet Island	Latin America and the Caribbean
Brazil	Latin America and the Caribbean
British Indian Ocean Territory	Sub-Saharan Africa
Brunei Darussalam	South-eastern Asia
Bulgaria	Eastern Europe
Burkina Faso	Sub-Saharan Africa
Burundi	Sub-Saharan Africa
Cabo Verde	Sub-Saharan Africa
Cambodia	South-eastern Asia
Cameroon	Sub-Saharan Africa
Canada	Northern America
Cayman Islands	Latin America and the Caribbean

Central African Republic	Sub-Saharan Africa
Chad	Sub-Saharan Africa
Chile	Latin America and the Caribbean
China	Eastern Asia
Christmas Island	Australia and New Zealand
Cocos (Keeling) Islands	Australia and New Zealand
Colombia	Latin America and the Caribbean
Comoros	Sub-Saharan Africa
Congo	Sub-Saharan Africa
Congo, Democratic Republic of the	Sub-Saharan Africa
Cook Islands	Other
Costa Rica	Latin America and the Caribbean
Cote d'Ivoire	Sub-Saharan Africa
Croatia	Southern Europe
Cuba	Latin America and the Caribbean
Curacao	Latin America and the Caribbean
Cyprus	Western Asia
Czechia	Eastern Europe
Denmark	Northern Europe
Djibouti	Sub-Saharan Africa
Dominica	Latin America and the Caribbean
Dominican Republic	Latin America and the Caribbean
Ecuador	Latin America and the Caribbean
Egypt	Northern Africa
El Salvador	Latin America and the Caribbean
Equatorial Guinea	Sub-Saharan Africa
Eritrea	Sub-Saharan Africa
Estonia	Northern Europe
Eswatini	Sub-Saharan Africa
Ethiopia	Sub-Saharan Africa
Falkland Islands (Malvinas)	Latin America and the Caribbean
Faroe Islands	Northern Europe
Fiji	Other
Finland	Northern Europe
France	Western Europe
French Guiana	Latin America and the Caribbean
French Other	Other
French Southern Territories	Sub-Saharan Africa
Gabon	Sub-Saharan Africa
Gambia	Sub-Saharan Africa
Georgia	Western Asia
Germany	Western Europe
Ghana	Sub-Saharan Africa
Gibraltar	Southern Europe
Greece	Southern Europe
Greenland	Northern America
Grenada	Latin America and the Caribbean

Guadeloupe	Latin America and the Caribbean
Guam	Other
Guatemala	Latin America and the Caribbean
Guernsey	Northern Europe
Guinea	Sub-Saharan Africa
Guinea-Bissau	Sub-Saharan Africa
Guyana	Latin America and the Caribbean
Haiti	Latin America and the Caribbean
Heard Island and McDonald Islands	Australia and New Zealand
Holy See	Southern Europe
Honduras	Latin America and the Caribbean
Hong Kong	Eastern Asia
Hungary	Eastern Europe
Iceland	Northern Europe
India	Southern Asia
Indonesia	South-eastern Asia
Iran (Islamic Republic of)	Southern Asia
Iraq	Western Asia
Ireland	Northern Europe
Isle of Man	Northern Europe
Israel	Western Asia
Italy	Southern Europe
Jamaica	Latin America and the Caribbean
Japan	Eastern Asia
Jersey	Northern Europe
Jordan	Western Asia
Kazakhstan	Other
Kenya	Sub-Saharan Africa
Kiribati	Other
Korea (Democratic People's Republic of)	Eastern Asia
Korea, Republic of	Eastern Asia
Kuwait	Western Asia
Kyrgyzstan	Other
Lao People's Democratic Republic	South-eastern Asia
Latvia	Northern Europe
Lebanon	Western Asia
Lesotho	Sub-Saharan Africa
Liberia	Sub-Saharan Africa
Libya	Northern Africa
Liechtenstein	Western Europe
Lithuania	Northern Europe
Luxembourg	Western Europe
Macao	Eastern Asia
Madagascar	Sub-Saharan Africa
Malawi	Sub-Saharan Africa
Malaysia	South-eastern Asia
Maldives	Southern Asia

Mali	Sub-Saharan Africa
Malta	Southern Europe
Marshall Islands	Other
Martinique	Latin America and the Caribbean
Mauritania	Sub-Saharan Africa
Mauritius	Sub-Saharan Africa
Mayotte	Sub-Saharan Africa
Mexico	Latin America and the Caribbean
Other (Federated States of)	Other
Moldova, Republic of	Eastern Europe
Monaco	Western Europe
Mongolia	Eastern Asia
Montenegro	Southern Europe
Montserrat	Latin America and the Caribbean
Morocco	Northern Africa
Mozambique	Sub-Saharan Africa
Myanmar	South-eastern Asia
Namibia	Sub-Saharan Africa
Nauru	Other
Nepal	Southern Asia
Netherlands	Western Europe
New Caledonia	Other
New Zealand	Australia and New Zealand
Nicaragua	Latin America and the Caribbean
Niger	Sub-Saharan Africa
Nigeria	Sub-Saharan Africa
Niue	Other
Norfolk Island	Australia and New Zealand
North Macedonia	Southern Europe
Northern Mariana Islands	Other
Norway	Northern Europe
Oman	Western Asia
Pakistan	Southern Asia
Palau	Other
Palestine, State of	Western Asia
Panama	Latin America and the Caribbean
Papua New Guinea	Other
Paraguay	Latin America and the Caribbean
Peru	Latin America and the Caribbean
Philippines	South-eastern Asia
Pitcairn	Other
Poland	Eastern Europe
Portugal	Southern Europe
Puerto Rico	Latin America and the Caribbean
Qatar	Western Asia
Reunion	Sub-Saharan Africa
Romania	Eastern Europe

Russian Federation	Eastern Europe
Rwanda	Sub-Saharan Africa
Saint Barthelemy	Latin America and the Caribbean
Saint Helena, Ascension and Tristan da Cunha	Sub-Saharan Africa
Saint Kitts and Nevis	Latin America and the Caribbean
Saint Lucia	Latin America and the Caribbean
Saint Martin (French part)	Latin America and the Caribbean
Saint Pierre and Miquelon	Northern America
Saint Vincent and the Grenadines	Latin America and the Caribbean
Samoa	Other
San Marino	Southern Europe
Sao Tome and Principe	Sub-Saharan Africa
Saudi Arabia	Western Asia
Senegal	Sub-Saharan Africa
Serbia	Southern Europe
Seychelles	Sub-Saharan Africa
Sierra Leone	Sub-Saharan Africa
Singapore	South-eastern Asia
Sint Maarten (Dutch part)	Latin America and the Caribbean
Slovakia	Eastern Europe
Slovenia	Southern Europe
Solomon Islands	Other
Somalia	Sub-Saharan Africa
South Africa	Sub-Saharan Africa
South Georgia and the South Sandwich Islands	Latin America and the Caribbean
South Sudan	Sub-Saharan Africa
Spain	Southern Europe
Sri Lanka	Southern Asia
Sudan	Northern Africa
Suriname	Latin America and the Caribbean
Svalbard and Jan Mayen	Northern Europe
Sweden	Northern Europe
Switzerland	Western Europe
Syrian Arab Republic	Western Asia
Taiwan, Province of China	Eastern Asia
Tajikistan	Other
Tanzania, United Republic of	Sub-Saharan Africa
Thailand	South-eastern Asia
Timor-Leste	South-eastern Asia
Togo	Sub-Saharan Africa
Tokelau	Other
Tonga	Other
Trinidad and Tobago	Latin America and the Caribbean
Tunisia	Northern Africa
Turkey	Western Asia
Turkmenistan	Other
Turks and Caicos Islands	Latin America and the Caribbean

Tuvalu	Other
Uganda	Sub-Saharan Africa
Ukraine	Eastern Europe
United Arab Emirates	Western Asia
United Kingdom of Great Britain and Northern Ireland	Northern Europe
United States of America	Northern America
United States Minor Outlying Islands	Other
Uruguay	Latin America and the Caribbean
Uzbekistan	Other
Vanuatu	Other
Venezuela (Bolivarian Republic of)	Latin America and the Caribbean
Viet Nam	South-eastern Asia
Virgin Islands (British)	Latin America and the Caribbean
Virgin Islands (U.S.)	Latin America and the Caribbean
Wallis and Futuna	Other
Western Sahara	Northern Africa
Yemen	Western Asia
Zambia	Sub-Saharan Africa
Zimbabwe	Sub-Saharan Africa

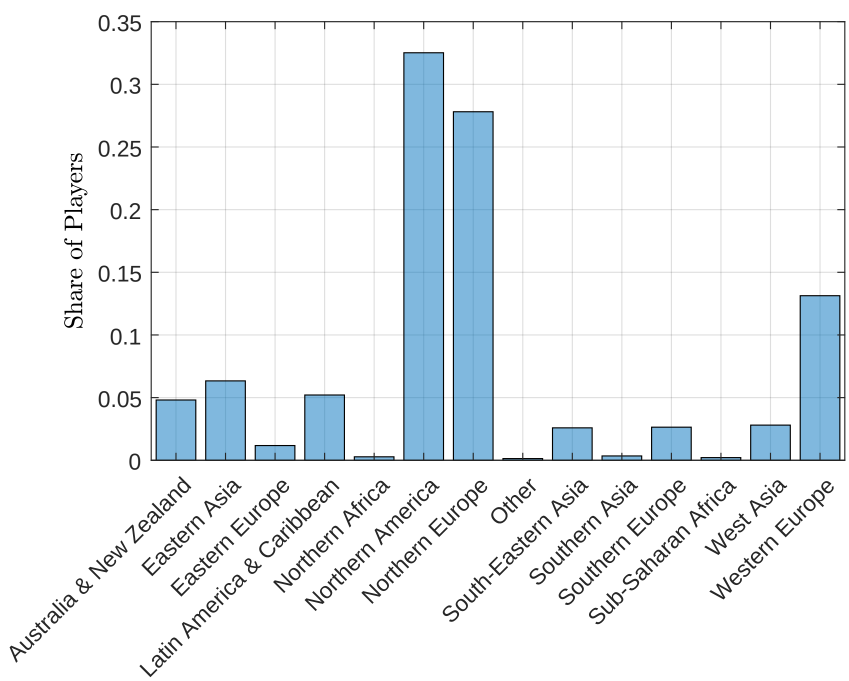

Figure A-26: Share of Players per Region

	1st gate at level 20		1st gate at level 40				Differences		
	Stars (1)		Stars (2)		No stars (3)		(1)-(2)	(1)-(3)	(2)-(3)
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.			
Avg. Snake Length	5,226	0,817	5,221	0,810	5,223	0,811	0,005	0,003	-0,002
Avg. Move Count	10,399	2,331	10,398	2,315	10,388	2,319	0,002	0,012	0,010
Avg. Final Score	24973,640	6554,295	24963,570	6544,456	24940,420	6570,346	10,070	33,220	23,150
Avg. Rounds per Level	3,932	5,665	3,989	5,787	3,934	5,703	-0,058	-0,002	0,055*
Number of Players	43,218		43,660		205,415				

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: This table provides evidence of balance between the different experimental groups. The sample includes all players in each group, as explained in section 3.2. It provides averages and associated t-tests (with unequal variances) across these groups. For the purposes of comparison, averages are calculated across the first 20 levels of the game, dropping rounds of the game spent on levels already previously cleared. The average snake length is the per-player average number of consecutive jellies assembled in a round of the game. The average move count counts instead the number different moves played in the game. The average final score reflects the aggregate performance of a player in a given round of the game. The average rounds per level is the number of attempts before first success.

Table A-15: Balance test: Comparison across groups