Pay to (Not) Play: Monetizing Impatience in Mobile Games

Taylor Lundy[1], Narun Raman[1], Hu Fu[2], Kevin Leyton-Brown[1]

[1] University of British Columbia
[2] Shanghai University of Finance and Economics, Key Laboratory of Interdisciplinary Research of Computation and Economics
tlundy@cs.ubc.ca, narunram@cs.ubc.ca, fuhu@mail.shufe.edu.cn, kevinlb@cs.ubc.ca

Abstract

Mobile gaming is a rapidly growing and incredibly profitable sector; having grown seven-fold over the past 10 years, it now grosses over $100 billion annually. This growth was due in large part to a shift in monetization strategies: rather than charging players an upfront cost ("pay-to-play"), games often request optional microtransactions throughout game-play ("free-to-play"). We focus on a common scenario in which games include wait times—gating either items or game progression—that players can pay to skip. Game designers typically say that they optimize for player happiness rather than revenue; however, prices for skips are typically set at levels that few players are willing to pay, leading to low purchase rates. Under a traditional analysis, it would seem that game designers fail at their stated goal if few players buy what they are selling. We argue that an alternate model can better explain this dynamic: players value tasks more highly as they are perceived to be more difficult. While skips can increase players' utilities by providing instant gratification, pricing skips too cheaply can *lower* players' utilities by decreasing the perceived amount of work needed to complete a task. We show that high revenue, high player utility, and low purchase rates can all coexist under this model, particularly under a realistic distribution of players having few buyers but a few big-spending "whales." We also investigate how a game designer should optimize prices under our model. An appendix of the paper with proofs, more comprehensive results and visualizations can be found at https://arxiv.org/abs/2312.10205.

Introduction

The gaming industry is larger than Hollywood and the music industry combined; currently valued at $245.10 billion USD, it is expected to reach $376.08 billion by 2028 (Intelligence 2023). While console and PC games have historically dominated the gaming category, mobile games have achieved seven-fold revenue growth since 2012 (Newzoo 2022) and accounted for a majority of gaming industry revenues in 2023 (Intelligence 2023). Reflecting this shifting marketplace, Microsoft recently announced plans to acquire Activision Blizzard for $68.7 billion USD (Microsoft 2022). Just two months after release, *Diablo Immortal*, Blizzard's franchise mobile title, surpassed $100 million USD; overall, mobile games

account for 40% ($1.08 billion USD) of Activision's revenue (Partis 2022; Activision 2022).

A key factor in the rise of mobile games is a shift in pricing models. Over 90% of mobile games are free to download ("free-to-play") and monetized by offering optional "micro-transactional" payments throughout game-play (Misiuk 2019; Seufert 2017; Tafradzhiyski 2023). While free-to-play is not a novel strategy (e.g., many platforms like Spotify offer free tiers), mobile games monetize a qualitatively different set of behaviours. For example, mobile games often allow users to spend money on virtual cosmetic items, extra lives that enable longer play sessions, or items that make the game easier. We are interested in understanding what player behavior causes these strategies and, in this paper, we focus on skips.

Games that monetize using skips have a gameplay loop requiring "grinding": completing tedious tasks, such as waiting for a timer to tick down. Players are offered the ability to skip the grind for a price. This style of monetization is so successful that an entire genre of games consisting solely of sequential timers and skips generated over $3 billion world-wide in 2021 (Pecorella 2013b; Moloco 2021).

While the success of skip-based monetization may be surprising, it is odder still how skips are priced. It is common to see talks at game design conferences discussing the perils of short-sighted monetization strategies (Greer 2013; Nussbaum and Telfer 2015; Krasilnikov 2020; Reilley 2023), or the importance of ensuring that players feel that purchases were "worth it" (Levy 2021; Engblom 2017; Reilley 2023). Successful game designers understand this well; when talking about fixing stagnating revenues in Clash of Clans, the chief game lead at the time, Eino Joas (2020) said, "retention always trumps monetization". Despite this, virtual goods such as skips are typically priced so high in practice that very few people are willing to buy them. On average, only 3.5% of players spend any money in-game and the majority of converted players, or 'minnows,' spend only $1 to $5 a month and contribute to less than 15% of the revenue (Brightman 2016; AppsFlyer 2020). The majority of revenue comes from 'whales,' who spend more than $25 per month on average and account for less than 15% of buyers (Shi, Xia, and Huang 2015). Under traditional economic models, high average utilities and prices few buyers can afford are incompatible. It might, therefore, seem either that designers are failing to make most players happy or that they truly care more about

revenue than player utility.

We show that existing game designs can indeed yield high player utility—along with high revenue—under an alternative model of player value. In this model, each individual player values skips because they dislike waiting. However, skips can decrease *all* players' values for virtual goods the more they are used. We believe that players value signaling that they achieved a high rank by grinding for it and cheap skips degrade this value by making progress purchasable. Modeling how the purchasing of skips degrades value requires modeling both how knowledgeable players are about other players' purchasing habits and how they incorporate this knowledge into their valuations. We make these modeling decisions based on the design dimensions of the game: the more immersive or social a game is the more sensitive players' values for playing are to different signals of grinding. We use this model to explain existing pricing strategies, but also show that, under some simplifying assumptions, game designers can leverage tools from traditional mechanism design to optimize prices.

We now describe what follows in the remainder of this paper. We start by discussing related work studying utility models grounded in difficulty, cost, or scarcity; offering evidence about ways in which mobile game players vary; and studying the monetization of mobile games. We then formally introduce our model of player utility where players have value functions that increase with the rarity of skip usage. We model three types of mobile game players, differing in their responsiveness to other players' behaviors: fully-sensitive, price-sensitive, and insensitive. With our model in hand we offer two main results. First, we demonstrate that, when players' value functions are price-sensitive, our model can explain the counter-intuitive pricing observed in practice, giving conditions under which high prices yield both good utility and good revenue. Furthermore, we show that these conditions are satisfied under player distributions with few buyers but a sufficient number of "whales." We illustrate these theoretical results in several examples and show that they also hold empirically in the more complicated fully-sensitive setting. Second, we prove that when players' value functions are insensitive, a simple pricing scheme approximates optimal revenue. We also show that this same pricing scheme approximates optimal revenue even when value functions are price-sensitive. We then test the simple pricing scheme with simulations, showing that in a variety of different settings, this scheme is competitive with other, natural pricing schemes. Finally, we conclude our work.

Related Work

A long line of literature—both in game theory and beyond—explores the idea that an object's value to an agent can be impacted by the difficulty or cost of obtaining it. For example, this idea is implicit in the idea of conspicuous consumption (Thorstein 1899; Battalio, Kagel, and Carl 1991), which posits that agents value goods in part because they can be used to signal wealth. In such models, higher prices and scarcity are primary drivers of value. Conspicuous consumption in mobile gaming is a well-documented phenomenon (Martin 2008; Lehdonvirta 2009; Cai, Wohn, and Freeman 2019). Directly related to our work, Geng and Chen (2019) investigated the freemium business model, finding that conspicuous consumption is a main driver of premium content pricing. In other cases, players can be motivated to signal their skill or commitment to a game rather than their wealth (Carter and Björk 2016). Further work has shown how social factors can exacerbate conspicuous consumption (O'Cass and McEwen 2004). Realizing this, mobile games often incorporate social networks to reinforce the conspicuous value of the virtual goods they offer (Fields and Cotton 2011; Choi and Kim 2004; Grieve et al. 2013). Goetz and Lu (2022) use data from online games to show that being matched with other players who have acquired virtual items makes a given player more likely to purchase them.

Of course, a player's perception of value also depends on intrinsic factors, such as their degree of impatience. In this vein, Hanner and Zarnekow (2015) showed how player purchasing behaviour can be predicted based on past purchases. Additionally, these intrinsic factors can lead to very different player behaviour. Game designers often report that in real-world player populations, the majority of revenue comes from a select group (around 2-3%) of "whale" players (Zatkin 2017; Shi, Xia, and Huang 2015; AppsFlyer 2020).

A key theme of this paper is that optimal pricing in mobile games must carefully trade off the value players assign to rewards that are difficult to obtain and the impatience that they experience as a result. The relationship between these two factors has been studied in various other contexts. Hsiao and Chen (2015); Evans (2016); Keogh and Richardson (2018) study how these factors affect the motivations behind purchases in mobile games. Recent work by Sepulveda, Besoain, and Barriga (2019) looks at how optimizing game difficulty can impact player satisfaction.

Player happiness is a first-order consideration in the monetization of mobile games: games that rely on microtransactions need players to keep coming back to play more and pay more. The link between retention and revenue has been well studied in both academia and the game industry, with many approaches for increasing retention being empirically validated and widely deployed (Nojima 2007; Pecorella 2013a). For example, Xue et al. (2017); Huang, Jasin, and Manchanda (2019) study how to adjust the difficulty of a game or opponents to keep players from leaving.

Payments for skipping past obstacles yield considerable revenue in mobile games (Pecorella 2013b). This style of difficulty-reducing monetization has started to gain attention in the optimization literature. Sheng et al. (2020) study the revenue and retention implications of offering items that make the game easier, focusing on games that provide these items in exchange for watching an ad. Further afield, Jiao, Tang, and Wang (2021) study the impact of selling skill-boosting items in multiplayer, competitive games.

Other recent papers study additional challenges related to game monetization that overlap less with our own work. Examples include optimizing pricing for loot boxes, which are lotteries over virtual items (Chen et al. 2021), and characterizing tradeoffs between perpetual and subscription-based licensing (Dierks and Seuken 2020).

Our Model

The most straightforward rationale for players valuing game playing is that the game mechanic itself is enjoyable: it is fun to blast meteors out of the sky; to explore an open world; to solve puzzles; etc. Many mobile games have mechanics that defy such an explanation. Instead, in order to complete a level or unlock an item, players must simply wait for a timer to tick down. It is hard to believe that there is anything fun about this kind of game mechanic (e.g., most people dislike waiting for a bus or a checkout clerk). Furthermore, players would not value the opportunity to skip timers if waiting were what made the game enjoyable.

Our model rests on a sharply different assumption: that players value achieving tasks because of the perceived difficulty of completing them. This means that as achievements become scarcer in the player population, their value increases. FarmVille crops that take days to grow have high value because they are hard to obtain; a Town Hall Level 12 in Clash of Clans is desirable because obtaining one requires 2 years and 2 days of waiting through timers (Crux 2023). This is similar to existing models of conspicuous consumption except that the cost is denominated in time rather than money.

While there is value in difficult gameplay, there is also frustration. We assume that a player's demand for skips comes from a time discount in the reward for completing the task, which we model as their type $\gamma \in [0, 1]$. We allow for players to differ in their impatience: they draw their types privately and i.i.d. from some prior continuous distribution F_γ. Furthermore, value in difficult gameplay can be shaped by the way other players in the player base accomplish their own tasks: players trying to signal their commitment to grinding care about how they compare to others. Skip prices can therefore affect the conspicuous value of a task, as lower skip prices decrease the fraction of players who accomplished a task by waiting. Typically, every player views the same price for a task and has no way to tell if another player completed a task by waiting or by skipping. Thus, the perceived value of a task falls with the cost of skipping it, even if the number of players who accomplish it is held constant.

How players' values for signalling grinding are impacted by skip prices depends heavily on how games are designed. We divide mobile games into three broad categories—midcore, casual, and hypercasual—and argue that each gives rise to a different level of sensitivity to other players' behavior. First, mid-core games (e.g., Clash of Clans, Game of War) incorporate in-game communities as a way to foster conspicuous value and immersion. For instance, a crucial part of game-playing in Clash of Clans is joining 'clans' with other players. Players can see the progress of other members in their clan and easily deduce if a clan member's rank was gained by purchasing or waiting. We say that players in these games are *fully-sensitive* and assume they know the full type distribution. We assume that fully-sensitive players determine their value based on both the price of a skip as well as the type distribution of other players.

Definition 1.1. *A fully-sensitive player with type γ drawn from type distribution F_γ, at the beginning of a task has projected value $\gamma v(F_\gamma, p)$ for finishing the task without skipping,*
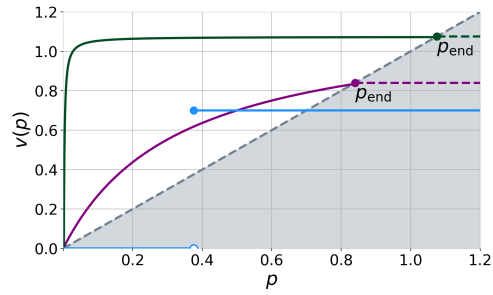

Figure 1: The purple and green lines are examples of value functions satisfying all of our assumptions for fully-sensitive and price-sensitive players, and the blue line satisfies our assumptions for insensitive players.

when the skip price is p.

We insist that these value functions maintain some nice properties with respect to price. First, the limit case of zero-cost skips is nonsensical (they would always be chosen), so we assume that $v(F_\gamma, 0) = 0$. We also assume that value is monotonically increasing and concave in p. These two conditions mean that as the price increases, the probability of sale for skips weakly decreases and the value of the accomplishment weakly increases. To avoid the uninteresting cases where players obtain non-positive utility at all skip prices we assume that the $v'(F_\gamma, 0)$ is greater than 1. Finally, we assume that there exists a point $p_{\text{end}} > 0$ where $v(F_\gamma, p_{\text{end}}) = p_{\text{end}}$ at which point the value function becomes a constant p_{end}. Otherwise, the designer can give players arbitrarily high utility by setting prices arbitrarily high. Figure 1 illustrates two value functions that satisfy the properties above, with the shaded region indicating values for $v(F_\gamma, p)$ that yield negative utility at price p.

Casual games (e.g., AdVenture Capitalist), on the other hand, are generally simple, easy to learn, and easy to play. The play style of casual games reflects their simplicity: most play for a short time while in a line at a supermarket or waiting for a friend (Turkowski 2023). Designers of casual games generally do not incorporate in-game social communities and often obscure information about other players' purchasing habits such as whether they paid to accomplish tasks. Still, players can be very invested in casual games, spending large amounts time and money to progress (Pecorella 2013b; AppsFlyer 2020). We assume these players do not know the type distribution induced by the given game but instead have a prior over what the type distribution might be. In this case, the true type distribution cannot affect player value, so a *price-sensitive* player's value function depends only on price.

Definition 1.2. *A price-sensitive player with type γ at the beginning of a task has projected value $\gamma v(p)$ for finishing the task without skipping, when the skip price is p.*

We assume the same properties about the relationship between value and price hold as just described.

Lastly, we look at players of hypercasual games (e.g., Cow Clicker, Cookie Clicker, Clicker Heroes). Such players play many hypercasual games and are not invested in any particu-

lar game (Turkowski 2023). We therefore consider it unlikely that players of hypercasual games are knowledgeable about the distribution of players for any given game. Furthermore, hypercasual players predominantly play to relieve stress, pass time between daily activities, or while multitasking, making it unlikely that they are reasoning about how price impacts their enjoyment (Facebook 2019). While skip prices may still need to be set high enough that the game does not feel pointless, beyond this threshold players value functions are *insensitive* to skip prices.

Definition 1.3. *An insensitive player with type γ at the beginning of a task has projected value $v_c(p) = \gamma v$ when the skip price $p > c$ for some constant c, and 0 otherwise, for finishing the task without skipping. We assume c is the same for all players.*

For ease of exposition, the rest of this section will define player utilities with respect to price-sensitive values. The corresponding definitions for fully-sensitive and insensitive values are very similar and can be found in Appendix A of the extended version.

Definition 1.4. *At the beginning of a task, the marginal value a player with type γ has for a skip at price p is $v_m(p) = (1 - \gamma)v(p)$. The utility a player receives if they purchase a skip at price p is $u(v(p), p) = v(p) - p$. We assume players are utility maximizers. Therefore, the utility a player receives in a task is $u_{\max}(v(p), p, \gamma) = \max(\gamma v(p), v(p) - p)$.*

Our technical results in this paper leverage a common distributional condition from the economics literature called a *monotone hazard rate* (MHR).

Definition 1.5. *A distribution $F(x)$ satisfies the MHR condition if the inverse hazard rate, $\frac{1 - F(x)}{f(x)}$, is monotone non-increasing.*

Optimal Pricing for Sensitive Players

We begin with studying price-sensitive players and show that they naturally give rise to the phenomenon of utility-optimal and revenue-optimal pricing coinciding at high prices. In particular, we show this when the game exhibits a *whale distribution:* (1) a large majority of the population never buys anything; and (2) a small fraction of players (whales) that generate a large majority of overall revenue. We then empirically demonstrate distributions where the results with price-sensitive players carry over to fully-sensitive players and others where their pricing behavior is different. Note that while we will be characterizing the distributions that give rise to high prices, we only allow the game designer to control skip prices. For the entirety of this section we focus on the setting with a single unrepeated task.

Price-Sensitive Players

The technical work of this section unfolds in three parts. First, we show a condition on the type distribution under which it is utility optimal to set skip prices arbitrarily high and that it is satisfied by distributions with feature (1). Second, we give a characterization of how the value function can drive the revenue-optimal price higher and demonstrate the impact of distributions with feature (2). Finally, we leverage these two conditions to construct distributions where the revenue optimal payment still yields high utility. In this section, for analytic convenience, we only consider value functions that are differentiable everywhere.

We first need a description of the total utility generated by skips in a task.

Proposition 2.1. *The expected utility (of all players) from a single task is $\mathbf{E}_\gamma[u_{\max}(\cdot)] = F_\gamma\left(1 - \frac{p}{v(p)}\right)(v(p) - p) + \left(1 - F_\gamma\left(1 - \frac{p}{v(p)}\right)\right)\mathbf{E}_\gamma[\gamma|\gamma > 1 - \frac{p}{v(p)}]v(p)$. We denote the utility-optimal price as $p_{\text{util}} = \arg\max_p \mathbf{E}_\gamma[u_{\max}(\cdot)]$, and the resulting utility it achieves as $\mathcal{U}_{\max}$.*

All proofs for this section can be found in the extended version.

We begin by giving a condition implying that the utility-maximizing price is arbitrarily close to p_{end}. The following condition is of interest on its own, as it characterizes some cases when offering skip microtransactions becomes detrimental to expected player happiness.

Theorem 2.2. *If $F_\gamma\left(\frac{v'(0) - 1}{v'(0)}\right) \leq \frac{v'(p_{\text{end}})}{1 - v'(p_{\text{end}})}\mathbf{E}_\gamma[\gamma]$, then $\forall \epsilon > 0$, $\mathbf{E}_\gamma[u_{\max}(v(p_{\text{end}}), p_{\text{end}}, \gamma)] \geq \mathcal{U}_{max} - \epsilon$.*

The proof for this theorem follows from rearranging the definition of optimal utility and can be found in the appendix. It is useful to interpret this condition through the lens of feature (1) of whale distributions. The left hand side of the condition represents the highest possible probability of sale across all non-zero prices. Meanwhile, the right hand side contains the expectation of γ which grows as the population becomes more patient. Whale distributions, given a fixed value function, make this condition easier to satisfy; as the number of non-paying players increases, the probability of sale on the left hand side decreases and the expected value on the right hand side increases. The intuition behind this result is quite simple: when there are enough patient players in the population, more total utility is gained by raising the price to make them happy than by providing cheap skips as an easy out for impatient players.

We now have a condition on the type distribution that would imply that the utility-optimal game design sets a price with zero probability of sale. The next step is to characterize distributions for which the revenue-optimal price monetizes only a small fraction of players. First, we describe the revenue generated by skips in a single task with respect to the distribution over impatience. It is useful to keep in mind that a player with high γ is a patient player, therefore the distribution over $1 - \gamma$ is a distribution over "impatience."

Proposition 2.3. *The revenue generated from offering a skip price p from a single task is $\mathcal{R}(v(p), p, \gamma) = p\left(1 - F_{1-\gamma}\left(\frac{p}{v(p)}\right)\right)$.*

Lemma 2.4. *The revenue-optimal price $p_{\text{rev}} \geq p^*$ if for all $p \leq p^*$, $\frac{1 - F_{1-\gamma}\left(\frac{p}{v(p)}\right)}{f_{1-\gamma}\left(\frac{p}{v(p)}\right)} \geq \frac{p}{v(p)}\left(1 - \frac{pv'(p)}{v(p)}\right)$. Furthermore, if $F_{1-\gamma}$ satisfies MHR then the optimal price is the point where this is satisfied with equality.*

Here, we interpret what this says for type distributions and their revenue-optimal payments. This lemma's first condition is a lower bound on the inverse hazard rate of the impatience distribution. The larger the domain over which this holds, the larger lower bound we have for the revenue-optimal payment. As the density of the most impatient players increases, the domain where this condition is satisfied can only grow. One might worry that having the high density of patient non-buyers necessary for Theorem 2.2 could violate Lemma 2.4. However, the density of players who would never buy are not considered directly by this condition; for large enough γ, $1 - \gamma \neq p/v(p)$ for any price, so their density never appears in the inverse hazard rate.

The second part of the lemma gives us an analogue to virtual values. Traditionally, virtual values are defined with respect to a constant value function. Suppose we define a constant value function as $v = v(p^*)$ where p^* is some fixed point, the resulting virtual values will be greater for any point $p^{**} > p^*$ than defined with respect to $v(p^{**})$. To see this, note that the virtual value for a constant value function would be identical except for the term $(1 - pv'(p)/v(p))$ which is always less than one while the value function is increasing. This demonstrates the upward pressure an increasing value function has on pricing.

Putting together the conditions from Theorem 2.2 and Lemma 2.4, we characterize a price that provides both high utility and high revenue.

Corollary 2.5. *Let p_{rev} be the revenue optimal price and m_{rev} be the slope of the value function at $v(p_{\mathrm{rev}})$. If the condition in Theorem 2.2 holds, then setting the revenue-optimal price p_{rev} achieves utility $\mathbf{E}_\gamma[u_{\max}(v(p_{\mathrm{rev}}), p_{\mathrm{rev}}, \gamma)] \geq \mathcal{U}_{\max} - m_{\mathrm{rev}}(p_{\mathrm{end}} - p_{\mathrm{rev}})$.*

This simple fact demonstrates that we can ensure near-optimal utility by setting prices high enough to reach a near-flat region of the value function. As mentioned in the beginning of the section, while high revenue-optimal and utility-optimal prices seem at odds with traditional mechanism design, many game designers are often encouraged to increase their prices for microtransactions and find increases in both overall retention and revenue (Levy 2016).

Fully-Sensitive Players

The results of the previous section depend on both the value function and the type distribution. It is not obvious if these results hold when the value function itself depends on the type distribution. We now show empirically that for value functions that are linear in the probability of sale we can still get high utility and revenue optimal prices but sometimes optimal prices can be pushed down if the value function grows too slowly with price. The fully-sensitive value functions that we consider satisfy the desired properties we listed, for specifics on how to construct such value functions and a more detailed breakdown of pricing behavior in this setting, see the extended version of this paper.

For a while, our intuition from the price-sensitive model holds; create enough patient players and the utility-optimal price goes to the max. However, if there are too many patient players in the population, the value function plateaus

(a) Truncated exponential distributions with rate parameter $\lambda \in \{1, 5, 15\}$.

(b) Utility optimal prices (▼, ▼, ▼) changing as density of patient players concentrates.

Figure 2: Utility-optimal price as type distributions concentrate on patient players. Value functions plotted on the right generated by corresponding type distributions on the left.

sharply and the trend flips: the utility-optimal price starts to drop (seen in Figure 2b). This is because the utility for non-buyers in this region of price-insensitivity remains relatively constant, while buyers' utility drops linearly with p.

We find that heavier tails usually mean higher revenue-optimal prices. However, it becomes more difficult to drive prices to the extremes in cases where the utility-optimal price is also high. This is because as the function gets flatter the upward pressure exerted on the price lessens. We offer more discussion and visualizations in the extended version.

This section shows that when value functions are sensitive to changes in price across the domain, the designer's job is relatively simple. The utility-optimal price often dominates the revenue-optimal price, so setting a high price and lowering it over time provides little risk in losing players. In fact, this is a strategy that game designers are sometimes advised to implement (Levy 2016). However, when value functions have large regions of insensitivity, the trade-off between buyer utility and revenue is non-obvious. This trade-off is the focus of the next section.

Optimal Pricing for Insensitive Players

We now tackle optimizing revenue for players insensitive to skip prices. We also note that we can transform fully- and price-sensitive value functions into an insensitive value function, allowing the results in this section to provide a good approximation when these sensitive value functions have large regions of low sensitivity. By definition, insensitive players have a minimum price below which their value is 0. However, this point does not affect the analysis, so we assume it is at 0 for the remainder of the section and let $v := v_c(p)$. For more discussion see Appendix D in the extended version.

In this section, we leverage the analytic tractability afforded when players are insensitive to investigate the harder problem of repeated task pricing. Our main result is a simple pricing scheme that gives a constant approximation to the revenue of the optimal pricing scheme. This approximation ratio holds even if the designer is allowed to set multiple prices based on how much of the task players complete. Not only is our simple pricing scheme interesting theoretically, it is also practically relevant: we often see pricing schemes in

which a single price is offered regardless of how much time has elapsed, even in mid-core (*Clash of Clans*) and casual games (*AdVenture Capitalist*). We end the section by testing our results in richer environments through simulations.

Repeated Tasks

We now introduce the further modeling assumptions and notation that we will need to study repeated tasks. We make two main assumptions: player's types remain constant across tasks, and players are myopic to future value. That is, they only consider the value of the task at hand rather than accounting for future value they might obtain by completing subsequent tasks. This matches experimental results on human behaviour showing that people often bracket their choices, optimizing piecemeal, rather than considering them jointly (Tversky and Kahneman 1985). Mobile games often exacerbate this effect, obscuring the amount of time needed to complete subsequent tasks. For example, Clicker Heroes only shows the timers for current tasks, hiding all future wait timers.

Definition 3.1. *At the beginning of each task k, the designer sets a take-it-or-leave-it price p_k. This results in an infinite sequence of prices $\{p_k\}_{k \in \mathbb{N}}$.*

One might imagine that game designers try to extract as much revenue from their players as possible. However, their task is not so simple. The mobile game marketplace is highly saturated with offerings in every conceivable permutation of aesthetic elements; e.g., *Cow Clicker*, *Cookie Clicker*, *Planet Clicker*, *Room Clicker*, *Battery Clicker*, and *Candy Clicker* are all real games! In such a marketplace, players who are retained in one round may not be retained in the next; designers have to worry about losing players to these outside options at every round of the game. We characterize the appeal of these outside options as a retention threshold resampled from a static distribution at the beginning of every task. This utility check means a game is always at a risk of losing its least satisfied players.

Definition 3.2. *At the end of each task $k \in [1, \infty)$, all players realize a retention threshold $r_k \sim F_r$, where $\text{supp}(F_r) = [0, \infty)$. A player quits the game immediately if $u_{\max}^{(k)}(v(p_k), p_k, \gamma) < r_k$, where $u_{\max}^{(k)}(v(p_k), p_k, \gamma) = \max(\gamma v(p_k), v(p_k) - p_k)$.*

An Upper Bound on Optimal Revenue

Characterizing the optimal sequence of prices is difficult, especially when one must worry about the type distribution changing over time due to retention. However, the revenue generated when players' types are known to the designer is not only easy to characterize but upper bounds the revenue when types are unknown. Within this *known types* setting, the optimal strategy sets $p_i = (1 - \gamma_i)v$ for each player i which is guaranteed to sell and static across tasks. We use the term *unknown types* for the setting where the designer has only a prior over players' types and must price each round anonymously. Note that in either setting the designer only has access to the distribution over retention thresholds. We assume designers exponentially discount future rewards

characterized by β. The proofs in this section can be found in the extended version.

Proposition 3.3. *The revenue generated from offering a static skip price p_i to a player i with known type γ_i and game designer discount β is $\mathcal{R}(v, p_i, \gamma_i) = \frac{p_i}{1 - \beta F_r(v - p_i)}$, if $p_i \leq (1 - \gamma_i)v$ and 0 otherwise.*

We now make an assumption on the retention distribution that is similar to regularity (many common distributions satisfy it e.g., uniform, exponential, Pareto).

Definition 3.4. *We say a retention distribution $F_r(x)$ is retention regular if $x - \frac{1 - \beta F_r(v - x)}{\beta f_r(v - x)}$ is monotone non-decreasing.*

Proposition 3.5. *The optimal price for a player i given known types, under retention regularity of the retention distribution, is $p_i = \min\left((1 - \gamma_i)v, \frac{1 - \beta F_r(v - p)}{\beta f_r(v - p)}\right)$, where p is the unique solution to the equation $q = \frac{1 - \beta F_r(v - q)}{\beta f_r(v - q)}$.*

Notice that retention adds a constant upper bound on the price we should charge. For players with high enough marginal value, charging a price that is too high may leave them with low utility, risking revenue from future tasks. Conversely, if a player has low marginal value for a skip, full surplus extraction may still leave them with high enough utility to be retained. We build on these observations for our simple pricing scheme later in this section.

Of course, designers are not restricted to offering only one skip per task, for example they could charge less if the player waits for half the task. We show that setting multiple prices can generate more revenue in Appendix E in the extended version. While even finding a closed form for such a pricing scheme is difficult, they are also upper bounded by the known types setting.

We note that without any distributional assumption on the marginal value distribution there is an unbounded revenue gap between known and unknown types, so we have no simple tool for bounding the performance of these complex pricing schemes.

Theorem 3.6. *For any constant c, there exists a retention distribution F_r and marginal value distribution F_m such that a game designer with known types achieves c times more revenue than the revenue optimal price for unknown types.*

This theorem relies on reducing our setting to a well known setting in the literature known as the equal-revenue distribution (Hartline and Roughgarden 2009; Hartline 2013). A natural question is whether there is a restricted class of distributions in which pricing schemes with unknown types can still perform competitively. Distributions with a monotone hazard rate (MHR) are one such example.

Simple Skip Pricing

The technical content unfolds in two steps: (1) we bound the gap in revenue at a given round between the known types setting and a simple pricing scheme using a result from the literature on distributions with a monotone hazard rate (MHR); and (2) we show our pricing scheme maintains at least as many players as the known types setting and always preserves

MHR. Combining these two facts guarantees a simple pricing scheme can $1/e$-approximate the optimal revenue. See the extended version for the definition of revenue in the unknown types setting.

We begin by defining how retention impacts the marginal value distribution across tasks in the unknown types setting. Note that at the start of the game, the distribution over players' marginal values is $F_m(p) = \Pr_\gamma[(1-\gamma)v \leq p]$.

Definition 3.7. *A marginal value distribution $F_m^{(k)}$ at task $k > 0$ if a retention threshold of r_{k-1} was drawn at the end of task $k - 1$ is defined as: $F_m^{(k)}(x) := F_m^{(k-1)}(x \mid u_{\max}(v(p), p_{k-1}, \gamma) \geq r_{k-1})$, where $F_m^{(0)}(x) := F_m(x)$.*

We formalize the observation of "threshold" pricing from the known types setting for the unknown types setting.

Definition 3.8 (Myerson Threshold (MT) Pricing). *First, the designer computes $p_r = {(1 - \beta F_r(v-p))}/{\beta f_r(v-p)}$. Then at each task k, they compute the Myerson optimal posted price with respect to the marginal value distribution at that task: $p_k^{\mathrm{myer}} = {(1 - F_m^{(k)}(x))}/{f_m^{(k)}(x)}$. Finally, for task k, the designer charges $p_k = \min(p_r, p_k^{\mathrm{myer}})$.*

Note that the p_r computation is the same as in Proposition 3.5. We can think of this as the retention distribution setting a constant upper bound on the price across all tasks.

We leverage the following result to lower-bound the revenue obtained by the Myerson optimal price.

Lemma 3.9. *(Hartline 2013, Theorem 4.37) For any marginal value distribution $F_m^{(k)}$ satisfying the MHR property, the expected value is at most e times the expected revenue generated in one round by setting the optimal posted price.*

Finally, we show that as long as the marginal value distribution satisfies MHR at the first task, threshold pricing achieves a $1/e$-approximation of the optimal revenue.

Theorem 3.10. *For every retention distribution F_r satisfying retention regularity and marginal value distribution F_m satisfying MHR, Myerson Threshold (MT) pricing generates a $1/e$ fraction of the revenue in the known types setting.*

Note that for any discount factor, marginal values decrease as the task goes on. Since in the known-type case the designer can guarantee a sale, they will always sell in the first round. This means the upper bound used in the proof of Theorem 3.10 applies equally to the multiple prices setting.

Corollary 3.11. *For every retention distribution F_r satisfying retention regularity and marginal value distribution F_m satisfying MHR, Myerson Threshold (MT) pricing generates a $1/e$ fraction of the revenue of any complex pricing scheme.*

Simulations

In this section, we simulate populations of players to study what features of the problem make MT pricing perform well. We also test the robustness of MT pricing by adding new ingredients that break away from some of our previous assumptions. Namely, we allow the population to grow from newly acquired players and players to be heterogeneous with respect to retention draws. Full details of the parameters used in the simulations and the formal descriptions of these new

ingredients can be found in Appendix G in the extended version. We note that even when adding these new ingredients, MT pricing gets at least a $1/e$ fraction of the optimal known types' revenue. Therefore, we compare MT pricing's performance to other simple pricing schemes to learn what features of the problem makes each option perform well. Histograms of the relative performance of MT pricing and the next best alternative for each of the following sections can be found in Appendix H in the extended version.

Comparison to Myerson and Retention Threshold Pricing
We begin by comparing two alternative pricing strategies: at each round, charge only the Myerson optimal price or the retention threshold price. The results suggest that the lower price often achieves the higher revenue. Setting a retention threshold price achieves higher revenue in over 20% of the instances and most of these instances are when Myerson takes a risk in setting a high price and loses too many players.

This matches the advice we see from the gaming community. Rather than losing too many players trying to make a quick buck, MT pricing keeps enough players around to monetize in the next task. In fact, the revenue generated from MT pricing is within 1% of the revenue from the max of the two other pricing schemes in 82% of cases and strictly outperforms the next best option in 12% of cases.

Viral Population Growth The gap in performance between MT pricing and the alternatives is small in most settings. This should not be surprising; MT charges the minimum of Myerson and the retention threshold and in many cases this results in exactly the same strategy as one of them, including retaining the same individuals.

Independently Drawn Retention Once again, the strategy that sets the lower price often generates higher average revenue. However, Myerson pricing now consistently loses players when it gambles on a high price, eliminating any extra gains it could have achieved.

Lowering Payments The main finding of our simulations thus far is that lower prices are usually better. One might wonder if the MT pricing algorithm could be improved by charging $p_k = \min(p_r, cp_k^{\mathrm{myer}})$ where c is a constant less than one. We re-examine a subset of previous settings that showed evidence of lower prices being performant and find that decreasing the scaling factor can be relatively safe for large enough values of c. However, for values lower than $1/2$, the variance in the performance not only increases but average performance decreases.

Conclusion

Mobile games give rise to various phenomena that are difficult to explain with standard economic models. This paper aims to explain the paradoxical relationship between utility and revenue when selling skips. We find that optimal pricing requires maintaining high player utility within and across tasks; high enough to have value in the game but low enough to monetize long-term. We believe that we have only scratched the surface of a rich economic domain; there are many other phenomena within mobile games (and beyond) upon which our behavioural model could shed light.

Acknowledgements

References

Activision. 2022. 2021 Annual Report - Activision Blizzard — Investor relations. https://investor.activision.com/static-files/d7b4f08d-213b-4bd5-a41b-7497baa9c106. Accessed: 2022-06-15.

AppsFlyer. 2020. 2020 Mobile Gaming Apps Report. https://liftoff.io/resources/report/2020-mobile-gaming-apps/. Accessed: 2022-01-26.

Battalio, R.; Kagel, J.; and Carl, K. 1991. Experimental Confirmation of the Existence of a Giffen Good. *The American Economic Review*, 81(4): 961–970.

Brightman, J. 2016. Only 3.5% of gamers make in-app purchases - appsflyer. https://www.gamesindustry.biz/only-3-5-percent-of-gamers-make-in-app-purchases-appsflyer. Accessed: 2022-01-26.

Cai, J.; Wohn, D. Y.; and Freeman, G. 2019. Who Purchases and Why? Explaining Motivations for In-Game Purchasing in the Online Survival Game Fortnite. *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, 391–396.

Carter, M.; and Björk, S. 2016. Cheating in candy crush saga. *Social, casual and mobile games: the changing gaming landscape*, 261–74.

Chen, N.; Elmachtoub, A. N.; Hamilton, M. L.; and Lei, X. 2021. Loot box pricing and design. *Management Science*, 67(8): 4809–4825.

Choi, D.; and Kim, J. 2004. Why people continue to play online games: In search of critical design factors to increase customer loyalty to online contents. *CyberPsychology & behavior*, 7(1): 11–24.

Crux, P. 2023. Clash of Clans Village Progress Tracker. https://pixelcrux.com/Clash_of_Clans/Village_Progress_Tracker/App. Accessed: 2023-09-02.

Dierks, L.; and Seuken, S. 2020. Revenue Maximization for Consumer Software: Subscription or Perpetual License? *arXiv preprint arXiv:2007.11331*.

Engblom, S. 2017. Balancing Cards in Clash Royale. https://www.youtube.com/watch?v=bHLQQh8Ctu4&t=969s. Accessed: 2022-01-26.

Evans, E. 2016. The economics of free: Freemium games, branding and the impatience economy. *Convergence*, 22(6): 563–580.

Facebook. 2019. Genre and great games report. https://www.facebook.com/fbgaminghome/marketers/build-great-games/game-design-hub/genre-great-games-report. Accessed:2023-07-14.

Fields, T.; and Cotton, B. 2011. *Social game design: Monetization methods and mechanics*. CRC Press.

Geng, W.; and Chen, Z. 2019. Optimal Pricing of Virtual Goods with Conspicuous Features in a Freemium Model. *International Journal of Electronic Commerce*, 23(3): 427–449.

Goetz, D.; and Lu, W. 2022. Peer Effects from Friends and Strangers: Evidence from Random Matchmaking in an Online Game. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 285–286.

Greer, E. 2013. Building games for the long-term: Maximizing monetization and player satisfaction. https://www.gdcvault.com/play/1018265/Building-Games-for-the-Long. Accessed: 2022-01-26.

Grieve, R.; Indian, M.; Witteveen, K.; Tolan, G. A.; and Marrington, J. 2013. Face-to-face or Facebook: Can social connectedness be derived online? *Computers in human behavior*, 29(3): 604–609.

Hanner, N.; and Zarnekow, R. 2015. Purchasing behavior in free to play games: Concepts and empirical validation. In *2015 48th Hawaii International Conference on System Sciences*, 3326–3335. IEEE.

Hartline, J. D. 2013. Mechanism design and approximation. *Book draft. October*, 145.

Hartline, J. D.; and Roughgarden, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, 225–234.

Hsiao, K.-L.; and Chen, C.-C. 2015. What drives in-app purchase intention for mobile games? An examination of perceived values and loyalty. *Electronic Commerce Research and Applications*.

Huang, Y.; Jasin, S.; and Manchanda, P. 2019. "Level Up": Leveraging skill and engagement to maximize player gameplay in online video games. *Information Systems Research*, 30(3): 927–947.

Intelligence, M. 2023. Global gaming market: 2023 - 2027: Industry share, size, growth - mordor intelligence. https://mordorintelligence.com/industry-reports/global-gaming-market. Accessed: 2023-08-06.

Jiao, Y.; Tang, C. S.; and Wang, J. 2021. Selling virtual items in free-to-play games: Transparent selling vs. opaque selling. *Service Science*, 13(2): 53–76.

Joas, E. 2020. Clash of Clans: Bigger, Better, Battle Pass. https://www.youtube.com/watch?v=jl9JCmaF9Y0. Accessed: 2022-01-26.

Keogh, B.; and Richardson, I. 2018. Waiting to play: The labour of background games. *European Journal of Cultural Studies*, 21(1): 13–25.

Krasilnikov, V. 2020. Monetization Design: The Dark Side of Gacha. https://www.youtube.com/watch?v=LnCOkQ-f8AQ&t=1316s. Accessed: 2022-01-26.

Lehdonvirta, V. 2009. Virtual Item Sales as a Revenue Model: Identifying Attributes that Drive Purchase Decisions. *Electronic Commerce Research*, 9(1): 97–113.

Levy, E. 2016. How to design in-game purchases - youtube. https://www.youtube.com/watch?v=vc2oNgRPTR0. Accessed: 2022-01-26.

Levy, T. 2021. Nature vs. Nurture: Unpacking Player Spending in F2P Games. https://www.youtube.com/watch?v=HvUn1XOURk4&t=693s. Accessed: 2022-01-26.

Martin, J. 2008. Consuming code: use-value, exchange-value, and the role of virtual goods in Second Life. *Journal For Virtual Worlds Research*, 1(2).

Microsoft. 2022. Microsoft to acquire Activision Blizzard to bring the joy and community of gaming to everyone, across every device. https://news.microsoft.com/2022/01/18/microsoft-to-acquire-activision-blizzard-to-bring-the-joy-and-community-of-gaming-to-everyone-across-every-device/. Accessed: 2022-02-06.

Misiuk, K. 2019. Free or Paid Mobile Games - What's More Profitable? https://indiewatch.net/2019/06/16/free-or-paid-mobile-games-whats-more-profitable/. Accessed: 2022-01-26.

Moloco. 2021. How big is the hyper casual games market? https://www.moloco.com/en/blog/hyper-casual-games-market-size. Accessed: 2022-02-06.

Newzoo. 2022. The games market and beyond in 2021: The Year in numbers. https://newzoo.com/insights/articles/the-games-market-in-2021-the-year-in-numbers-esports-cloud-gaming/. Accessed: 2022-02-06.

Nojima, M. 2007. Pricing models and Motivations for MMO play. *Digital Games Research Association*.

Nussbaum, S.; and Telfer, A. 2015. In it for the Long Hall: How Wooga boosts long-term retention. https://www.gdcvault.com/play/1022070/In-It-for-the-Long. Accessed: 2022-01-26.

O'Cass, A.; and McEwen, H. 2004. Exploring consumer status and conspicuous consumption. *Journal of Consumer Behaviour*, 4(1): 25–39.

Partis, D. 2022. Diablo immortal generates $100m in lifetime revenue. https://www.gamesindustry.biz/diablo-immortal-generates-100m-in-lifetime-revenue#:~:text=Diablo%20Immortal%20has%20surpassed%20%24100,two%20months%20after%20its%20release. Accessed: 2022-06-15.

Pecorella, A. 2013a. Building Games for the Long Term: Pragmatic F2P Guild Design. Accessed: 2022-02-01.

Pecorella, A. 2013b. The Rise and Rise of Idle Games. Accessed: 2022-01-04.

Reilley, D. 2023. The evolution of Game Monetization, retention and acquisition (presented by Amazon prime gaming, Voodoo, Nifty Games, Publicis Groupe and Scuti / Lenderwize). https://www.gdcvault.com/play/1029433/The-Evolution-of-Game-Monetization. Accessed: 2023-10-06.

Sepulveda, G. K.; Besoain, F.; and Barriga, N. A. 2019. Exploring Dynamic Difficulty Adjustment in Video Games. In *2019 IEEE CHILEAN Conference on Electrical, Electronics Engineering, Information and Communication Technologies (CHILECON)*, 1–6. IEEE.

Seufert, E. B. 2017. Apps Aren't Dying, but not Everyone Knows How to Make Money on Mobile: Mobile Dev Memo by Eric Seufert. https://mobiledevmemo.com/apps-arent-dying-making-money-on-mobile/. Accessed: 2022-01-26.

Sheng, L.; Ryan, C. T.; Nagarajan, M.; Cheng, Y.; and Tong, C. 2020. Incentivized actions in freemium games. *Manufacturing & Service Operations Management*.

Shi, S. W.; Xia, M.; and Huang, Y. 2015. From Minnow to Whales: An Empirical Study of Purchase Behavior in Freemium Social Games. *International Journal of Electronic Commerce*, 2015: 20–2.

Tafradzhiyski, N. 2023. In-app purchases. https://www.businessofapps.com/guide/in-app-purchases/#:~:text=In%2Dapp%20purchases%20account%20for,37.8%25%20from%20paid%20app%20downloads. Accessed: 2022-01-26.

Thorstein, V. 1899. The Theory of the Leisure Class. *New York: Mentor Books*.

Turkowski, J. 2023. Types of gamers: Casual and hardcore gamers as target groups. https://instreamly.com/posts/types-of-gamers-casual-and-hardcore-gamers/. Accessed: 2023-03-26.

Tversky, A.; and Kahneman, D. 1985. The framing of decisions and the psychology of choice. In *Behavioral decision making*, 25–41. Springer.

Xue, S.; Wu, M.; Kolen, J.; Aghdaie, N.; and Zaman, K. A. 2017. Dynamic difficulty adjustment for maximized engagement in digital games. In *Proceedings of the 26th International Conference on World Wide Web Companion*, 465–471.

Zatkin, G. 2017. Awesome Video Game Data 2017. Accessed: 2022-01-18.