

Loot Box Pricing and Design

Ningyuan Chen

Rotman School of Management, University of Toronto, Toronto, ON M5S3E6 nychen@ust.hk

Adam N. Elmachtoub

Department of Industrial Engineering and Operations Research & Data Science Institute, Columbia University, New York, NY 10027, adam@ieor.columbia.edu

Michael L. Hamilton

Katz Graduate School of Business, University of Pittsburgh, Pittsburgh, PA 15260, mhamilton@katz.pitt.edu

Xiao Lei

Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, xl2625@columbia.edu

In the online video game industry, a significant portion of the revenue is generated from microtransactions, where a small amount of real-world currency is exchanged for virtual items to be used in the game. One popular way to conduct microtransactions is via a *loot box*, which is a random allocation of virtual items whose contents are not revealed until after purchase. In this work, we consider how to optimally price and design loot boxes from the perspective of a revenue-maximizing video game company, and analyze customer surplus under such selling strategies. Our paper provides the first formal treatment of loot boxes, with the aim to provide customers, companies, and regulatory bodies with insights into this popular selling strategy.

We consider two types of loot boxes: a traditional one where customers can receive (unwanted) duplicates, and a unique one where customers are guaranteed to never receive duplicates. We show that as the number of virtual items grows large, the unique box strategy is asymptotically optimal among all possible strategies, while the traditional box strategy only garners 36.7% of the optimal revenue. On the other hand, the unique box strategy leaves almost zero customer surplus, while the traditional box strategy leaves positive surplus. Further, when designing traditional and unique loot boxes, we show it is asymptotically optimal to allocate the items uniformly, even when the item valuation distributions are heterogeneous. We also show that when the seller purposely misrepresents the allocation probabilities, their revenue may increase significantly and thus strict regulation is needed. Finally, we show that even if the seller allows customers to salvage unwanted items, then the customer surplus can only increase by at most 1.4%.

Key words: loot boxes; e-commerce; video games; bundling; probabilistic selling

1. Introduction

With the recent explosion of online and mobile gaming over the last decade (Perez 2018), the idea of games-as-a-service (GaaS) has been widely accepted as a way to provide video game content as a continuing revenue stream. Recently, the video game industry standard

has shifted towards the *freemium* model, where access to a game is freely given to customers, and in-game virtual items can be acquired via microtransactions. In other words, virtual items that help players in the game are purchased with in-game or real-world currency. In many of these games, microtransactions are conducted via a randomized mechanism known in the video game industry as a *loot box*. A loot box is a random allocation of virtual items, the contents of which are revealed after the purchase is complete. While the concept of a loot box is not new – for instance, a pack of baseball cards is a form of loot box – modern versions of loot boxes like in Fig. 1 have proliferated in online video games on mobile, console, and PC platforms over the last decade. In online games such as *Dota 2*, *FIFA 20*, *PlayerUnknown’s Battlegrounds*, and many others, loot box sales are a core source of revenue. In these games, customers purchase loot boxes which contain a random subset of virtual items such as character costumes, cosmetic upgrades, players, virtual cards, etc. In 2018 alone, more than \$30 billion dollars in sales were conducted via loot boxes (Wright 2018).

Figure 1 Loot Box Example.



Note. Depicted is a loot box in the popular video game *Dota 2*, where it is also referred to as a “treasure”. The customer may purchase the loot box for \$2.49, after which they receive one of the five items depicted on the screen.

Despite its popularity in the gaming industry, the use of loot boxes has invited controversy and criticism recently (Tassi 2018), with several development platforms enacting specific regulations in response (Apple March-8, Google July-23). For instance, there have been issues regarding the transparency of the contents and probability of outcomes from loot boxes (Fingas 2018) which have led to regulatory investigations (Tassi 2017). Another

issue is that loot boxes have been associated with gambling, both by the media (Webb 2017) and academia (Drummond and Sauer 2018, Zendle and Cairns 2018). This led the United States Congress to introduce a bill outlining new loot box regulations (Romm and Timberg 2019) and in August 2019 the Federal Trade Commission held a workshop on the matter (Holt 2019), which included the presentation of an early version of this paper (Sinclair 2019).

In spite of this negative publicity, loot box selling remains as popular as ever (Batchelor 2017). In order to properly address the issues of transparency and gambling via regulation, we believe that it is fundamental to understand the economic motivation behind loot box selling. Why do video game companies prefer such a business strategy? How does it compare to traditional strategies such as bundling or separate selling? What are the behavioral considerations motivating customers to keep purchasing loot boxes? A rigorous framework for the operations of loot boxes would provide valuable insights for customers, companies, and regulators, which is precisely the focus of the work.

There are several salient features that distinguish the loot box mechanism and virtual items from other traditional industries. First, the virtual items have zero marginal cost, can be copied infinitely by the seller, and have no value outside the game. Second, customers engage in repeated interactions with the seller, and may potentially buy a large number of items *sequentially*. Third, the seller is fully aware of a customer's current collection, and has control over the loot box allocation rules. These unique features render the models designed for other types of products inadequate and call for a new revenue management framework to specifically analyze loot box strategies.

In this work, we provide a model to analyze the optimal pricing and design of loot boxes for revenue-maximizing sellers. The model incorporates two types of commonly used loot box strategies. A *unique loot box* allocates items to customers that they do not currently own. A *traditional loot box* allocates items randomly to customers regardless of whether they already own a copy, which may result in obtaining unwanted duplicates. In both strategies, the seller may control the price, number of items allocated, allocation probabilities, and salvage value of the loot boxes. To understand the advantages and disadvantages of loot box strategies, we compare them to two traditional selling mechanisms: *separate selling*, where every virtual item is sold separately for a known price, and *grand bundle selling*, where customers pay a one-time fixed amount for access to all virtual items.

Next we provide a summary of our contributions and findings.

1. We propose the first mathematical modeling framework for selling loot boxes. Customers are endowed with i.i.d. private valuations for all items, and sequentially purchase loot boxes until their expected utility becomes negative. We show that almost no dominance relations hold between any of the four selling strategies we consider.
2. Motivated by the fact that the number of virtual items in a video game is often in the hundreds or thousands, we turn our attention to the setting where the number of items offered tends to infinity. In this asymptotic regime, both the grand bundle and unique box strategies are optimal, while the traditional box strategy only generates 36.7% as much revenue. Surprisingly, we show that the expected number of purchases from the traditional box strategy is roughly the same as the unique box strategy, although the optimal price of the traditional box is lower. However, the unique box strategy provides no customer surplus in this regime, in contrast to the traditional box strategy which leaves a positive customer surplus. We then connect our asymptotic results back to practice by conducting a numerical study to show that our findings still hold when the catalog size is finite and moderately sized.
3. Next, to accommodate the scenario that items may belong to classes with different rarities and values in the game, we consider the case where the valuation distributions are heterogeneous across items. When the number of items is large, we show that the optimal allocation probabilities of a loot box is a random draw among all available items, *independent of* how customers value items in any class. Thus under an optimal price and allocation rule, a rare class is less likely to be drawn because there are less items in the class, not because customers value it highly. We also show that the seller may gain significant revenue if they can successfully deceive consumers into believing a false set of allocation probabilities, even if such allocation probabilities are accurate in expectation. We conclude that regulation is needed to protect consumers from such a practice.
4. Finally, we consider an additional design aspect where customers are allowed to salvage, or return, unwanted items. We show that traditional box strategies may earn more revenue with salvage systems and are guaranteed to dominate separate selling. Surprisingly however, we show that introducing a salvage system in a traditional box strategy can only increase customer surplus by at most 1.4%. On the other hand, as

unique box strategies are already optimal, salvage systems cannot increase their revenue. However such systems do allow sellers to trade-off between revenue and customer surplus in a smooth fashion.

1.1. Literature Review

While loot box selling has not been previously studied in the revenue management literature to the best of our knowledge, our work draws inspiration from and is related to several areas across operations management, computer science, and economics.

In the operations management literature our work connects with the dual streams of papers on opaque selling and bundle selling. Loot boxes are an example of opaque selling, which is the practice of selling items where some features of the item are hidden from the customer until after purchase. Recent works (Jiang 2007, Fay and Xie 2008, Elmachtoub et al. 2020, Elmachtoub and Hamilton 2020) have focused on opaque selling as a tool to manage imbalanced customer demand or induce opportunities for price discrimination. Our loot box framework diverges from the standard opaque selling models in a number of key ways. First, we consider the performance of loot boxes in isolation, as opposed to many models in which the opaque option is sold in conjunction/in competition with traditional sales channels (Shapiro and Shi (2008), Jerath et al. (2010), Chen et al. (2014), Huang et al. (2017))). Second, we model complex, repeated interactions between the loot box seller and a customer interested in obtaining a catalog of items as opposed to prior work which has focused on customers who want at most one item. Third, in our loot box model we do not have finite inventories, which diverges from the literature on using opaque products to balance inventory (Gallego and Phillips (2004), Gallego et al. (2004), Xiao and Chen (2014), Elmachtoub et al. (2019)).

Our work also resembles and references the work on bundling. As in the bundling literature, the loot box is a way to sell products to markets of customers with demand for many items. In particular, we compare our loot box selling mechanisms explicitly with the grand bundle mechanisms studied in the seminal work of Bakos and Brynjolfsson (1999), who show that pure bundling extracts almost all of the consumer surplus asymptotically. By leveraging results from the theory of random walks, we show that the unique box strategy (allocating one item at a time without replacement) can achieve a similar revenue to grand bundle selling without forcing the customer to choose between purchasing the whole catalog or nothing. This property allows loot boxes to circumvent many of the issues that

plague bundle selling in practice (c.f. Section 5.2 for a detailed discussion). From a technical stand point, this is a stark divergence from the typical techniques in the bundling literature which rely on concentration results to induce a single purchase of the entire catalog of items.

Mixed bundle strategies, strategies that allow customers to purchase the items from a menu offering both the grand bundle and the items individually, and randomized strategies, strategies that offer lotteries over possible allocations of items, have been considered recently in a stream of work on optimal mechanisms for selling items to additive buyers (Babaioff et al. (2020), Hart et al. (2014), Hart and Reny (2015), Briest et al. (2014), Abdallah et al. (2019), Abdallah (2018)). A loot box can be thought of as a particularly simple type of randomized bundle strategy where only a single lottery over all the items is offered, and from which a single item is allocated. A similar type of mechanism is considered in Briest and Röglin (2010). They study the optimal pricing of menus of *unit-demand bundles* under stylized valuation assumptions, while we study repeated interactions with a customer which is the main driver of loot box revenues. Moreover, the focus of their work is on the computational hardness of computing the prices for menus of such bundles. In contrast, our loot box mechanism is dynamic and computationally simple. Ma and Simchi-Levi (2016) considers bundle selling with return options in the presence of production costs, while our paper considers loot boxes with return options and no production costs.

Closest to our paper, in the sense that an individual customer dynamically purchases multiple items from the seller, is the work of Ferreira and Goh (2020). There the authors consider whether or not to offer the products in sequence or all at once, but do not consider any form of randomized selling strategies such as a loot box. The focus of their paper is on understanding the value of concealment in the context of fast fashion, whereas our loot box model does not conceal any part of the catalog, and the focus is on the choice of randomized strategy. There has also been a line of work where a customer makes decisions in multiple stages when faced with an assortment from the seller, although only at most one unit is purchased (Wang and Sahin 2017, Gallego et al. 2020, Golrezaei et al. 2018).

Finally, our work contributes to the emerging literature on operations management in video games. Chen et al. (2017) and Huang et al. (2018) investigate the problem of maximizing a player's engagement in video games. Sheng et al. (2020) considers the problem of incentivizing actions in freemium games. Jiao et al. (2020) considers whether the seller

should disclose an opponent’s skill level when selling in-game items that can increase the win rate. Our work is the first to investigate the popular practice of loot box selling via mathematical modeling.

2. Model and Preliminaries

We consider a revenue-maximizing monopolist selling a catalog of N distinct, non-perishable, virtual items. A random customer’s valuation for the items are described by non-negative i.i.d. random variables $\{V_i\}_{i=1}^N$, where each V_i is drawn from a distribution F . The mean and variance of V_i are denoted by μ and σ^2 , respectively, and are assumed to be finite. The assumption of i.i.d. valuations is reasonable when the items are cosmetic (such as character skins and customizations) or when items are of similar importance, both of which are common in many games that deploy loot boxes. In Section 4.2, we extend our model to address the case where items are vertically differentiated and can naturally be categorized into multiple classes based on their values or rarities.

We suppose that each customer is aware of all available items in the seller’s catalog as well as their own realized valuations for the items v_i for $i \in [N]$, where $[N]$ is used to represent the index set $\{1, \dots, N\}$. Each newly obtained item i gives the customer a one-time utility of v_i , which can be thought of as the lifetime value of the item in the game, and is assumed to be independent of the period in which it is received. Further, no customer values having duplicates of an item, meaning a customer’s valuation for a second unit of each item i is 0. For example, a character skin or cosmetic upgrade for the player’s avatar, once obtained, can be enjoyed for as long as the player engages with the game, and a second copy offers no additional value to the player. In some games, the seller provides a salvage mechanism through which the customer can obtain value from duplicate items by trading them in for (in-game or real-world) currency. We discuss this extension in Section 4.4.

A loot box can be formally defined as a random allocation of a single item to the customer, chosen according to a probability distribution over all N items. We note that the probability distribution is decided by the seller, and may or may not depend on the customer’s current inventory. There are also cases where multiple items are allocated in one loot box, which is an extension we consider in Section 4.1. Moreover, we assume that the customer always knows the actual allocation probabilities (that is, the probabilities of

receiving each item). This is consistent with industry practice, as sellers are often forced to announce the allocation probabilities, either by government issued customer protection regulations (Tassi 2017) or by edict of the games distributor (Apple March-8, Google July-23). In Section 4.3, we consider extensions where the seller may misrepresent the allocation probabilities.

We now describe the sequence of events in our loot box model, which capture a single customer repeatedly interacting with the seller. Before the arrival of the customer, the seller announces the price and allocation probabilities of the loot box. We consider each purchasing event to be a discrete period and emphasize that periods do not necessarily correspond to any particular unit of time, i.e. periods can be thought of as occurring in rapid succession (for a player eager to complete their collection) or occurring with long gaps between purchases (for a more judicious player). In each period t , we let $S_t \subset [N]$ denote the index set of distinct items that the customer owns before opening the loot box in period t . Thus, $S_1 = \emptyset$. Based on the price, allocation probabilities, and the customer's private valuations for items in $[N] \setminus S_t$, the customer decides whether or not to purchase the loot box. We assume customers are utility-maximizing and will purchase if their expected utility from purchasing is non-negative, otherwise the customer will not purchase further loot boxes (they may however continue to play the game). We discuss the customer behavior in greater detail in Section 2.1.

We now formally describe the two forms of loot box selling that we focus on as well as two benchmark strategies known as grand bundle selling and separate selling.

1) Unique Box (UB): In the *unique box* strategy, the seller offers a loot box for a fixed price p in each period, with the guarantee that each purchase yields a new item that the customer does not yet own. Formally, the probability of receiving an item is 0 if $i \in S_t$, and $\frac{1}{|[N] \setminus S_t|}$ for $i \in [N] \setminus S_t$, i.e., uniform over all the items not currently owned by the customer. Fig. 2a shows an example of a unique box in practice. We let $\mathcal{R}_{UB}(p)$ be the normalized revenue of a unique box strategy that uses price p , i.e.,

$$\mathcal{R}_{UB}(p) := \frac{p \times \mathbb{E}[\# \text{ of Unique Box Purchases}]}{N}$$

and let $\mathcal{R}_{UB} := \max_p \mathcal{R}_{UB}(p)$.


2) Traditional Box (TB): In the *traditional box* strategy, the seller offers a loot box for a fixed price p in each period, with the guarantee that each purchase yields an item selected

uniformly at random from $[N]$, regardless of what the customer owns in S_t . Traditional boxes lead to the possibility of duplicate items during a customer’s purchasing process. Fig. 2b shows an example of a traditional box in practice. We let $\mathcal{R}_{TB}(p)$ be the normalized revenue of a traditional box strategy that uses a fixed price p , i.e.,

$$\mathcal{R}_{TB}(p) := \frac{p \times \mathbb{E}[\# \text{ of Traditional Box Purchases}]}{N}$$

and let $\mathcal{R}_{TB} := \max_p \mathcal{R}_{TB}(p)$.

Figure 2 Loot Boxes in Online Games.



(a) Unique Box

(b) Traditional Box

Note. The left panel shows a unique loot box in the popular online game *Dota 2*. The red square highlights that it is a unique box as the loot box always allocates a unique item. The right panel shows an implementation of a traditional loot box in the online game *PlayerUnknown’s Battlegrounds*.

We emphasize that, at first glance, it is not clear which loot box strategy generates more revenue. Intuitively, customers have higher valuations for unique boxes since they are guaranteed not to receive duplicates (recall we assume the utility derived from the second copy of any item is 0), which allows sellers to charge higher prices. On the other hand, although the seller may have to charge lower prices for traditional boxes, the selling volume may end up being higher because customers need to make more purchases in order to obtain new items. Indeed, for finite N , we provide instances where either strategy may dominate the other in Table 1.

Further, in both loot box strategies we assume the allocation probabilities are uniform over the remaining items/all items. We note that such allocation rules may not be optimal even though valuations for items are i.i.d. (see Example EC.1). In Section 4.2, we provide results on the asymptotic optimality of uniform allocation rules.

We shall compare and contrast these loot box models against two classic selling models: grand bundle selling and separate selling.

3) Grand Bundle (GB): In the *grand bundle* strategy, the seller offers a single bundle containing *all* N items for price Np . Customers no longer make dynamic decisions when a grand bundle is offered, but rather just make a single decision to purchase or not. The normalized revenue of a grand bundle strategy with price Np is

$$\mathcal{R}_{GB}(p) := \frac{Np\mathbb{P}\left(\sum_{i=1}^N V_i \geq Np\right)}{N},$$

and the optimal normalized revenue is denoted by $\mathcal{R}_{GB} := \max_p \mathcal{R}_{GB}(p)$. Fig. 3a shows an example of a grand bundle in practice.

4) Separate Selling (SS): In the *separate selling* strategy, the seller offers all items individually at the same price p . Since we assume the valuations V_i are i.i.d, the normalized revenue of a separate selling strategy with price p is

$$\mathcal{R}_{SS}(p) := \frac{Np\mathbb{P}(V_i \geq p)}{N},$$

and the optimal normalized revenue is denoted by $\mathcal{R}_{SS} := \max_p \mathcal{R}_{SS}(p)$. Fig. 3b shows an example of separate selling with uniform prices in practice.

Figure 3 Traditional Selling Strategies in Online Games.



(a) Grand Bundle



(b) Separate Selling

Note. The left panel shows an implementation of grand bundle selling in the online game *Brawlhalla*. All items can be unlocked for a one-time payment of \$19.99 via the *All Legends Pack*. The right panel shows an implementation of separate selling in the online game *Arena of Valor*. In this game each item (character) can be individually unlocked with uniform prices.

While there are many ways to sell virtual items, we restrict our attention to these four as we believe they capture the spirit of almost all strategies observed in practice.

Among them, the grand bundle and separate selling strategies provide two important benchmarks. Separate selling is the most common selling strategy in e-commerce. Grand bundle selling is also common for digital goods such as music and television, and has been shown to be able to fully extract the maximum possible revenue when N tends to infinity (Bakos and Brynjolfsson 1999). However, unlike the other three strategies, under the grand bundle strategy a customer must commit to purchasing *all* of the items or *none* of the items. In practice, this large upfront financial commitment may impair the grand bundle's performance. In contrast, separate selling and loot box strategies are more friendly to customers that prefer smaller purchases or have a smaller budget. Although the grand bundle selling may be prohibitive in practice, it serves as a useful theoretical benchmark since it is asymptotically revenue-optimal (see Section 5.2 for a detailed discussion).

2.1. Customer Behavior

We next describe how customers value loot boxes and make purchase decisions. We assume that customers are risk-neutral, and their valuation of a loot box is simply the expectation, over the allocation probabilities, of the valuation of the random item they will receive. Let U_t be the expected utility of opening a loot box in period t for a price p . Since the allocation probabilities are uniform, U_t has the following form:

$$\text{(Unique Box)} \quad U_t = \frac{\sum_{i \in [N] \setminus S_t} v_i}{N - |S_t|} - p, \quad \text{(Traditional Box)} \quad U_t = \frac{\sum_{i \in [N] \setminus S_t} v_i}{N} - p.$$

Naturally, to maximize the expected utility, customers would purchase the t^{th} loot box if $U_t \geq 0$. However, it is sometimes rational for a customer to purchase even if $U_t < 0$ for the prospect of higher utilities in future periods. The following example demonstrates that myopic behavior (purchasing if and only if $U_t \geq 0$) is not necessarily optimal for the customer.

EXAMPLE 1. Let $N = 2$ and consider the unique box strategy. Let the price of each loot box be $p = 1.6$. Consider a customer whose valuations of the two products are $(v_1, v_2) = (1, 2)$. If the customer is myopic, then they will not buy a single unique box since the expected utility from the first loot box purchase is $\frac{1+2}{2} - 1.6 < 0$. However, it can be shown by enumeration that the following purchasing strategy is optimal: buy a loot box in the first period. If the obtained item is product 2 where $v_2 = 2$, then stop purchasing. Otherwise purchase a second loot box, which is guaranteed to contain product 2. The expected net utility of this strategy is $\frac{1}{2}(2 - 1.6) + \frac{1}{2}(1 + 2 - 1.6 \times 2) = 0.1 > 0$. Thus, behaving myopically is strictly worse than the optimal strategy. \square

In Example EC.2 we generalize Example 1 and show that not only is the myopic strategy sub-optimal, but so is any policy that considers only a finite number of future states. Thus, in each period t , a perfectly rational customer needs to solve a high-dimensional and complex optimal stopping problem to decide whether or not to purchase. However, we believe it is both impractical and unrealistic for customers to find the the optimal purchasing strategy, as the state space of the corresponding optimal stopping problem increases exponentially in the number of items. Instead, we make the natural modeling assumption that customers are indeed myopic, i.e., they purchase if and only if their expected net utility for the next loot box U_t is non-negative. Theorem 1 shows that myopic behavior is asymptotically optimal for a customer facing unique box selling as the catalog of items grows large. Moreover, we show in Theorem 1 that myopic behavior is always optimal for a customer facing traditional boxes, lending additional support to our myopic assumption for customers facing loot boxes.

THEOREM 1 (Myopic Purchasing Behavior is Asymptotically Optimal).

- a) *For unique box selling, the average net utility under the myopic strategy converges to the average net utility of the optimal strategy as $N \rightarrow \infty$.*
- b) *For traditional box selling, the myopic purchasing strategy is optimal for all customers.*

Due to the apparent complexity and impracticality of computing the customer’s optimal purchasing policy and the fact that a myopic purchasing rule is near-optimal, we believe restricting to myopic purchasing behavior does not degrade the power of our models. Further, we note that in cases where the optimal purchasing strategy differs from the myopic strategy (i.e. unique boxes), the customer purchases strictly *more* loot boxes under the optimal behavior. To see this, note that a myopic customer will stop purchasing as soon a loot box gives negative utility, while an optimal customer may continue purchasing due to a positive expected future reward. Thus the revenue of a loot box strategy under the assumption of myopic behavior is a lower bound on the revenue when customers purchase optimally.

For the remainder of this paper, we shall assume customer behavior is myopic.

2.2. Comparing the Strategies for Finite N

In this section, we aim to understand the relations between the four strategies when N is finite. Specifically, we would like to establish dominance relations between the optimal

revenues of (UB), (TB), (SS), and (GB). In Proposition 1, we show that the normalized revenue of any of the strategies is always at most μ , and that the unique box strategy can never exceed the revenue of a grand bundle strategy.

PROPOSITION 1. *For any N and valuation distribution F , the following statements hold:*

- (a) $\mu \geq \max\{\mathcal{R}_{GB}, \mathcal{R}_{SS}, \mathcal{R}_{UB}, \mathcal{R}_{TB}\}$, *i.e., μ is a global upper bound on the normalized revenue.*
- (b) $\mathcal{R}_{GB} \geq \mathcal{R}_{UB}$, *i.e., grand bundle selling weakly dominates unique box selling.*

Proof. (a) For any strategy, the customer only makes a purchase when their expected utility is non-negative. Thus, the expected customer surplus is always non-negative. On the other hand, the total normalized expected welfare is always at most $\mathbb{E}[\sum_{i \in [N]} V_i]/N = \mu$. Together, these facts imply that the normalized revenue for any strategy is at most μ .

(b) Let p^* be the optimal price of the unique box strategy. An upper bound on \mathcal{R}_{UB} is then Np^* . For a customer to purchase the very first unique box, we must have that $\frac{\sum_i v_i}{N} \geq p^*$. Under the same condition, the customer would buy the grand bundle at price Np^* since $\sum_i v_i \geq Np^*$. Similarly, if $\sum_i v_i < Np^*$, the customer would not buy the first unique box nor the grand bundle at price Np^* . Therefore, the revenue from an optimal grand bundle strategy is at least as much as the optimal unique box strategy. \square

Unfortunately, outside of Proposition 1, there does not exist any other dominance relationships among the four selling strategies. In particular, the same argument in Proposition 1(b) does not extend to the comparison of grand bundle selling and traditional box selling. Although the condition for purchasing the first traditional box remains the same, a customer may end up buying strictly more than N boxes overall due to the possibility of duplicates. In Table 1 we give simple examples for which all of the 11 remaining possible relationships between the four selling strategies occur.

From Table 1, we can see that it is impossible to theoretically compare the selling strategies when N is small without imposing significant additional assumptions. Therefore, in the rest of the paper we focus on an asymptotic analysis where the number of items N tends to infinity. Fortunately, this is also well-motivated in the gaming industry where N , the number of items sold in a video game, is often in the hundreds or thousands. For example, in the popular online games *Dota 2* and *Overwatch*, the number of cosmetic items sold through loot boxes exceeds 3500. As we shall see, dominance relations among the four strategies naturally emerge in the asymptotic regime.

Table 1 Possible Relations Between (UB) , (TB) , (GB) and (SS) .

Relations	N	Valuation Distribution
$GB > UB, GB > TB, GB > SS$	3	$\mathbb{P}(V_i = 0.98) = 1/2, \mathbb{P}(V_i = 2.02) = 1/6, \mathbb{P}(V_i = 3.01) = 1/3$
$UB > SS, UB > TB, SS > TB$	10	$\mathbb{P}(V_i = 1) = 0.8, \mathbb{P}(V_i = 5) = 0.2$
$TB > UB, TB > SS$	3	$\mathbb{P}(V_i = 1.01) = 1/2, \mathbb{P}(V_i = 1.98) = 1/6, \mathbb{P}(V_i = 3.03) = 1/3$
$SS > GB, TB > GB$	2	$\mathbb{P}(V_i = 1) = \mathbb{P}(V_i = 100) = 1/2$
$SS > UB$	4	$\mathbb{P}(V_i = 1) = 3/10, \mathbb{P}(V_i = 10) = 7/10$

3. Asymptotic Analysis of Loot Box Strategies

In this section, we study the revenue of optimally priced loot box strategies as the number of items N in the catalog tends to infinity. Given the incomparability of the various selling strategies shown in Section 2.2 and the fact that N is often quite large in practice, this asymptotic analysis is quite natural. In this asymptotic regime, we shall show that an optimal unique box strategy earns a normalized revenue of μ per item (c.f. Theorem 2), whereas the optimal traditional box strategy earns a normalized revenue of only $\frac{\mu}{e} \approx 0.367\mu$ (c.f. Theorem 3). Since the expected normalized revenue of any selling strategy cannot exceed the mean valuation μ by Proposition 1, this result proves that unique box and traditional box strategies are asymptotically optimal and sub-optimal, respectively. Additionally, we can directly compare the performance of these two loot box strategies with grand bundle selling and separate selling in this regime. Using the strong law of large numbers, it is well known that the grand bundle also obtains a normalized revenue of μ (see Bakos and Brynjolfsson (1999) for a detailed discussion). On the other hand, the revenue of separate selling strategies depends explicitly on the distribution of customer valuations, and can earn anywhere between 0% and 100% of the normalized revenue.

THEOREM 2 (Asymptotic Revenue and Convergence Rate of UB). *The unique box strategy is guaranteed to earn*

$$\mathcal{R}_{UB} \geq \mu(1 - N^{-1/5}) \left(1 - \left(1 + \frac{2\sigma^2}{\mu^2} \right) N^{-1/5} - \frac{\sigma^2}{\mu^2} N^{-3/5} - \frac{\sigma^4}{\mu^4} N^{-4/5} - \left(\frac{\sigma^2}{\mu^2} + \frac{\sigma^4}{\mu^4} \right) N^{-6/5} \right).$$

Moreover,

$$\lim_{N \rightarrow \infty} \mathcal{R}_{UB} = \mu.$$

The proof of Theorem 2 follows by modeling the customer behavior dynamically as a random walk and explicitly constructing a sequence of prices that lead to the lower bound on the revenue. Specifically, we consider a random walk that captures the total utility a customer collects in each period and bound the number of purchases made. Unfortunately, the time a customer stops purchasing is not a stopping time since it depends on all of the customer's valuations, which are not known to the seller. Thankfully we are able to approximate the time that the customer stops purchasing by a true stopping time, and then leverage standard machinery to bound the total number of purchases. Finally, we show that setting a price of $p = \mu(1 - \frac{1}{N^{1/3}})$ leads to the desired result.

THEOREM 3 (Asymptotic Revenue and Convergence Rate of TB). *The traditional box strategy is guaranteed to earn*

$$\frac{\mu}{e} \log \left(\frac{1}{\frac{1}{e} + \frac{1 + \frac{\sigma^2}{\mu^2}}{N}} \right) \leq \mathcal{R}_{TB} \leq \frac{\mu}{e^{1-\zeta_N}(1 - N^{-\frac{1}{3}})} + \frac{(1 - N^{-\frac{1}{3}})\sigma^2 \log N}{\mu N^{\frac{1}{3}}},$$

where $\zeta_N = \sum_{i=1}^N \frac{1}{i} - \log(N) - \gamma$, and γ is the Euler-Mascheroni constant. Moreover,

$$\lim_{N \rightarrow \infty} \mathcal{R}_{TB} = \frac{\mu}{e}.$$

To prove Theorem 3, we construct a ‘backwards’ random walk that captures the total valuation collected by a customer, starting from the very last item they would purchase. We show that in our constructed random walk, the number of unique items collected corresponds to a stopping time, and again leverage standard machinery to bound the stopping time. Since duplicates are allowed, we must also account for the number of purchases required to collect a unique item, which depends on the number of items collected so far.

Theorems 2 and 3 highlight an important design aspect of loot boxes: the ability to monitor a customer's current inventory and appropriately control the allocations. With full information of a customer's inventory, a seller can implement unique boxes which are asymptotically revenue-optimal. Without this information, the seller is restricted to traditional loot boxes which garner only $\frac{1}{e}$ fraction of the optimal revenue. Note that this is not a lower bound, but rather an exact asymptotic limit of traditional loot box selling revenue. Moreover, both of these results hold for *any* underlying valuation distribution (with bounded first and second moments).

Next, we investigate the properties of optimal unique and traditional box strategies and study the limiting optimal price, sales volume, and customer surplus. We present these results in Theorem 4 and note that it provides several seemingly counter-intuitive insights into the differences between unique and traditional box strategies. For instance, one would expect that since unique boxes never leave customers empty-handed, they are preferred by customers. In addition, one may also expect that since traditional boxes yield duplicates, customers may tend to buy strictly more of them than unique boxes. Surprisingly, when optimally priced both of these intuitions are false.

THEOREM 4 (Insights into Loot Box Strategies).

- (a) *For the unique box strategy, as $N \rightarrow \infty$, the optimal price converges to μ . Further, the expected fraction of unique items collected by the customer converges to 1, the expected normalized number of loot boxes purchased converges to 1, and the expected normalized customer surplus converges to 0.*
- (b) *For the traditional box strategy, as $N \rightarrow \infty$, the optimal price converges to $\frac{\mu}{e}$. Further, the expected fraction of unique items collected converges to $1 - \frac{1}{e}$, the expected normalized number loot boxes purchased converges to 1, and the expected normalized customer surplus converges to $(1 - \frac{2}{e})\mu$.*

Theorem 4(a) states that the optimal price for unique boxes as N tends to infinity is approximately μ , and that customers purchase approximately the entire catalog. Since their average valuation is μ , this leaves them with no consumer surplus. On the other hand, Theorem 4(b) states that the optimal price for traditional boxes as N tends to infinity is approximately $\frac{\mu}{e}$, and that customers purchase approximately N boxes obtaining $1 - \frac{1}{e}$ fraction of the catalog of items. Although the consumers do not acquire the entire catalog, they are left with a positive normalized consumer surplus of approximately $(1 - \frac{2}{e})\mu$. To summarize, under both strategies the customer purchases approximately N loot boxes. However, the price is lower for traditional boxes, resulting in less revenue for the seller and more surplus for the consumer. Surprisingly, consumers are therefore better off with traditional boxes when duplicate allocations are allowed, since unique boxes lead to higher prices and vanishing customer surplus.

In light of Theorem 3, it is worthwhile to discuss why traditional boxes are popular among sellers in practice, given their substantially lower expected revenue in our model.

We posit three possible explanations. First, traditional boxes have existed well before the digital age in the form of Gachapon or as packs of Pokèmon cards, and the practice may continue as a hold over from those times. Second, there is long-term value in making sure consumers are left with positive surplus (c.f. Theorem 4), which is not explicitly included as an objective of the seller in our model. Third, the presence of a salvage system (resale market) may increase the revenue of a traditional box strategy since loot boxes can be sold at a higher price. We study this idea in detail in Section 4.4.

4. Loot Box Design

In this section, we extend the results of the previous section to handle various practical considerations and design aspects beyond the choice of unique versus traditional boxes. Recall that our basic model assumes that each loot box allocates one random item, that the valuations for all items are i.i.d., that the allocation distribution is uniformly random, and that customers obtain no value from duplicate items. In practice, these assumptions may sometimes be violated and thus we address them here. In Section 4.1, we discuss an extension where each loot box allocates multiple items. In Section 4.2, we discuss the case where there are multiple classes of items, and characterize the optimal allocation probabilities for potentially vertically differentiated items. In Section 4.3, we explore the regulatory concern when the seller deviates from the announced allocation probabilities. Finally, in Section 4.4, we consider the situation where the seller offers a salvage system (resale market), in which unwanted items can be salvaged by the customer for some return value.

4.1. Multi-item Loot Boxes

Although many games allocate one item at a time from their loot boxes, it is also a common practice to allocate multiple items at once (see Fig. EC.1 for an example). A classic example of loot boxes containing multiple items are Pokèmon or Baseball cards, which are sold in packs of ten or twelve. In practice, sellers may use a size- j loot box when the mean valuation of a single item μ is very low (e.g., less than \$0.10). In this case, selling multiple items in one box allows the seller to set a higher price, which helps reduce the number of transactions for the customer and allows prices to conform to market norms (e.g., the phenomenon of pricing at \$0.99).

In this section, we show that Theorems 2 and 3 can be extended to the case where loot boxes are of fixed size $j > 1$. We use \mathcal{R}_{UB}^j and \mathcal{R}_{TB}^j to denote the optimal revenue of the

size- j unique box and traditional box strategies, respectively. Proposition 2 shows that when j is fixed and N tends to infinity, the unique box strategy is still optimal and the traditional box strategy still only earns 36.7% as much.

PROPOSITION 2 (Multi-Item Loot Boxes). *For size- j loot boxes,*

$$\lim_{N \rightarrow \infty} \mathcal{R}_{UB}^j = \mu \quad \text{and} \quad \lim_{N \rightarrow \infty} \mathcal{R}_{TB}^j = \frac{\mu}{e}.$$

The proof uses a coupling argument. Specifically, we show that in the last period that a customer is willing to purchase a size- j box, they would have purchased a size-1 box as well. This reduces the total number of items purchased to the case studied in the previous section. In the case of the traditional box strategy, the seller needs to slightly decrease the price of the traditional size- j box as its value is now lower on average than a size-1 box, due to the possibility of more duplicates in a single box.

4.2. Optimizing Allocation Probabilities for Multiple Classes of Items

In previous sections, we assumed that valuations for all items are i.i.d., and that each item (*unowned* item in the case of unique boxes) was equally likely to be allocated by the loot box. In practice, these assumptions may not always hold. Often in online games, the items are explicitly grouped based on rarity or effectiveness. For instance, in the popular online game *PlayerUnknown's Battlegrounds*, customers may receive Mythic, Legendary, Epic, or Rare items from a loot box (see Fig. EC.2 for an example).

To model this phenomenon, suppose there are M different classes of items, and that each item $i \in [N]$ belongs to a specific class $m \in [M]$. Denote $G_m \subset [N]$ as the index set of items in class m and denote β_m as the proportion of the items belonging to class m , i.e., $\beta_m := |G_m|/N$ and $\sum_{m \in [M]} \beta_m = 1$. For each class of items m , the valuations for the items in that class are sampled i.i.d. from distribution F_m . We denote the mean and standard deviation of the valuations for items in class m by μ_m and σ_m . Let $\bar{\mu} := \sum_{m=1}^M \beta_m \mu_m$ be the expected valuation of a random item. For different classes, the distribution F_m may vary arbitrarily.

For asymptotic results, we shall suppose the number of items in each class grows proportionally with N , i.e., there are $\beta_m N$ items belonging to class m as N increases. The introduction of multiple item classes allows for some items to be significantly more valuable than others, and thus it is reasonable to consider non-uniform allocation probabilities

of the items. A loot box strategy is now characterized by a price p and the allocation probabilities of each item, which may depend on its class. Our goal is to characterize the revenue-optimal combination of price and allocation probabilities for loot boxes over multiple classes of items.

For unique box strategies, the optimal allocation probabilities may be non-uniform, dynamic, and depend on the current set of items owned by the customer. It is difficult to explain such policies to customers, let alone characterize the optimal allocation probabilities. Thankfully, for unique boxes there is a simple allocation and pricing strategy that is asymptotically optimal for the seller. Proposition 3 shows that a unique box strategy that simply allocates all unowned items uniformly at random, *completely ignoring the class*, is asymptotically optimal. Thus, for the t^{th} unique box, each unowned item is allocated with probability $\frac{1}{N-(t-1)}$.

PROPOSITION 3 ((UB) with Uniform Allocation is Asymptotically Optimal).

Suppose unique boxes allocate items uniformly at random, independent of class. Then we have that

$$\lim_{N \rightarrow \infty} \mathcal{R}_{UB} = \bar{\mu}.$$

Surprisingly, the proof of Proposition 3 is almost identical to that of Theorem 2. Although the distribution of the unowned items is no longer i.i.d., Wald's identity and Chebyshev's inequality continue to hold where appropriate.

Next, we shift focus to the allocation problem for the traditional box strategy. Once again, we shall show that a simple allocation policy is asymptotically optimal. Specifically, we show that it is optimal to allocate each item uniformly at random (w.p. $\frac{1}{N}$), *independent of the class*. To simplify the problem, we restrict our attention to class level allocation probabilities, i.e., allocation rules where all items in the same class have the same allocation probabilities but may differ between classes. We emphasize that class level allocation rules are common in practice (see Fig. EC.2 for examples). For a class level allocation rule, we let d_m be the probability of drawing an item in class m , and $\mathbf{d} = (d_1, \dots, d_M)$ be a probability vector. Thus, the probability of getting an item in class m is $\frac{d_m}{\beta_m N}$. A uniform allocation rule corresponds to the case where $\mathbf{d} = \beta$.

Let $Q^N(p, \mathbf{d})$ denote the normalized number of loot boxes purchased by a customer (i.e. $\frac{\mathbb{E}[\# \text{ Purchases}]}{N}$) and let $\mathcal{R}_{TB}(p, \mathbf{d})$ be the normalized expected revenue, both of which are a

function of the price p and class allocation probabilities \mathbf{d} . In Proposition 4 we show that for a carefully constructed price as a function of \mathbf{d} , the limit of $Q^N(p, \mathbf{d})$ can be characterized simply.

PROPOSITION 4. *Suppose a traditional box strategy follows a multi-class allocation rule \mathbf{d} and price $p = \sum_{m=1}^M d_m \mu_m e^{-\frac{d_m}{\beta_m} k}$ for some $0 < k < 1$. Then*

$$\lim_{N \rightarrow \infty} \mathbb{E} [Q^N(p, \mathbf{d})] = k.$$

To build intuition for Proposition 4, consider the case of a single class of items, and note from the proof of Theorem 4 that the number of unowned items after opening kN traditional boxes is roughly Ne^{-k} . If a customer stops after kN purchases, then their valuation for the next box is roughly μe^{-k} , meaning that a price of μe^{-k} will induce purchase up to that point but no further. The proposed price in Proposition 4 generalizes this intuition to the multi-class case. For example, when $d_m = \beta_m$ for all m and $k = 0$, the induced price in Proposition 4 is $p = \bar{\mu}$. Proposition 4 implies that the limiting normalized selling volume for this price with the uniform allocation strategy is 0. This agrees with our intuition, as the customer valuation drops below $\bar{\mu}$ after opening ϵN boxes for any fixed $\epsilon > 0$. Armed with Proposition 4, we find in Theorem 5 that as N tends to infinity, the simple strategy of setting the price to be $\frac{\bar{\mu}}{e}$ and the allocation probability vector to be β results in asymptotically optimal normalized revenue of $\frac{\bar{\mu}}{e}$.

THEOREM 5 ((TB) with Uniform Allocations are Asymptotically Optimal).

For traditional box with M classes, we have

$$\lim_{N \rightarrow \infty} \max_{p, \mathbf{d}} \mathcal{R}_{TB}(p, \mathbf{d}) = \lim_{N \rightarrow \infty} \mathcal{R}_{TB} \left(\frac{\bar{\mu}}{e}, (\beta_1, \dots, \beta_M) \right) = \frac{\bar{\mu}}{e}.$$

Proof. Consider a class allocation probability vector $\mathbf{d} = (d_1, \dots, d_M)$. For $p > \sum_{m=1}^M d_m \mu_m$, by the law of large numbers, the normalized selling volume will tend to 0. Thus we can focus on the case $p \leq \sum_{m=1}^M d_m \mu_m$. Note $\sum_{m=1}^M d_m \mu_m e^{-\frac{d_m}{\beta_m} k}$ equals $\bar{\mu}$ when $k = 0$, and decreases monotonically to 0 as $k \rightarrow \infty$. Thus for any price p , there exist a unique positive k such that $p = \sum_{m=1}^M d_m \mu_m e^{-\frac{d_m}{\beta_m} k}$. Recall by Proposition 4, if $p = \sum_{m=1}^M d_m \mu_m e^{-\frac{d_m}{\beta_m} k}$, $1 > k > 0$, then

$$\lim_{N \rightarrow \infty} \mathbb{E} [Q^N(p, \mathbf{d})] = k.$$

Using this identity we can write the limiting revenue function in terms of k , i.e.,

$$\lim_{N \rightarrow \infty} \mathcal{R}_{TB}(p, \mathbf{d}) = \lim_{N \rightarrow \infty} p \cdot \mathbb{E} [Q^N(p, \mathbf{d})] = k \sum_{m=1}^M d_m \mu_m e^{-\frac{d_m}{\beta_m} k} := \sum_{m=1}^M G_m(k).$$

Consider the m^{th} term of the revenue function, $G_m(k) = \mu_m d_m k e^{-\frac{d_m}{\beta_m} k}$. This function obtains its maximum at $k = \beta_m/d_m$, and the maximum value is $\beta_m \mu_m/e$, which is independent from the value of d_m . Hence, the limiting optimal revenue is bounded above by $\sum_{m=1}^M \beta_m \mu_m/e = \bar{\mu}/e$. For any \mathbf{d} , we can reach the the upper bound $\bar{\mu}/e$ if every component function reaches the maximum simultaneously in the limit, i.e., $k = \beta_m/d_m$ for all m . Since $\sum_{m=1}^M \beta_m = \sum_{m=1}^M d_m = 1$, the only possible limiting allocation is $d_m = \beta_m$, which is the uniform allocation. In this case, $k = 1$, and the corresponding price is $p = \sum_{m=1}^M d_m \mu_m e^{-\frac{d_m}{\beta_m} k} = \bar{\mu}/e$. Hence, the solution $p = \bar{\mu}/e$ with uniform allocation is asymptotically optimal with corresponding revenue $\bar{\mu}/e$. \square

Note that the allocation $\mathbf{d} = (\beta_1, \dots, \beta_M)$ is simply the uniform allocation over all the items (not over all classes). Thus Theorem 5 is a natural generalization of Theorem 3 to the multi-class case. The uniform allocation strategy with price $\bar{\mu}/e$ is asymptotically optimal, for any fixed number of classes and any set of valuation distributions. In this sense, Theorem 5 makes the decisions of a seller who adopts traditional box selling simple: instead of designing complicated allocation structures, just use uniform allocations and focus on the price. Further, Theorem 5 extends the asymptotic dominance of unique box strategies over traditional boxes to the case of multiple item classes; varying the allocation probabilities cannot close the gap in revenue between the two strategies.

We emphasize that for uniform allocations to be optimal in Theorem 5, the price for the loot box must be optimally chosen. For prices other than $\bar{\mu}/e$, uniform allocations may not be optimal. For example, suppose we have two classes with $\mu_1 = \$10$, $\mu_2 = \$5$ and $\beta_1 = \beta_2 = 0.5$, then the price $\frac{\bar{\mu}}{e} = \$2.76$ with uniform allocation is asymptotically optimal. However, if the seller uses the price $\$3$, then by optimizing over \mathbf{d} , we find that the asymptotically optimal class allocation probabilities are $\mathbf{d} = (0.514, 0.486)$.

Finally, we note that another natural selling mechanism to consider in this setting is to sell different classes separately via loot boxes of various prices. However, this mechanism is more complex and may create a sense of unfairness whereby wealthy players are able to obtain the high-value items more easily. Therefore, we believe that the simple allocation rule proposed in this section is preferable, in addition to being asymptotically optimal.

4.3. Transparency of the Allocation Probabilities

In previous sections, we assumed that both the customer and the seller believe and act according to the announced allocation probabilities. In practice, the seller sometimes may lie about the allocation probabilities by purposely using an allocation strategy different than the announced strategy. In this section, we discuss the potential implications of such deceit.

Consider a situation where the seller deviates from the posted allocation rule, but the customers still believe the posted allocation rule. As an example, such a situation occurred in the game *Monster Taming*, where the seller claimed the chance to receive a rare item was 1% whereas the actual odds were 0.0005% (Fingas 2018). In such cases, the duped customer may end up buying many more loot boxes due to the false announcement. In Example 2, we demonstrate that sellers can greatly increase their revenue by misrepresenting the allocation probabilities, and further, can do so in a way that is difficult to detect (unlike in the case of *Monster Taming*). In particular, Example 2 shows that a so-called *random perturbation strategy* can increase the revenue of a traditional loot box while adhering to the announced allocation rule in expectation, making such a deception hard to detect.

EXAMPLE 2. Consider a traditional box with a single class of items and price μ/e . Suppose the seller claims that a uniform allocation is used, but instead, the seller randomly chooses half of the items to be allocated with probability $\frac{1+\epsilon}{N}$ and the other half to be allocated with probability $\frac{1-\epsilon}{N}$. By Proposition 4, the normalized selling volume under a truly uniform allocation is asymptotically equal to 1. On the other hand, when the random perturbation is used, the traditional box can be regarded as a two-class traditional box, with $\mu_1 = \mu_2 = \mu$, and $\mathbf{d} = (0.5(1 + \epsilon), 0.5(1 - \epsilon))$. If the customer has complete information, then by Proposition 4, the normalized selling volume k is given by solving

$$\mu e^{-1} = \mu \left(\frac{1+\epsilon}{2} e^{-(1+\epsilon)k} + \frac{1-\epsilon}{2} e^{-(1-\epsilon)k} \right). \quad (1)$$

However, if the customer assumes the allocation is uniform, then the weight of two classes changes from $(\frac{1+\epsilon}{2}, \frac{1-\epsilon}{2})$ to $(\frac{1}{2}, \frac{1}{2})$ while the exponential terms in Eq. (1) remain the same. In this case k is given by solving

$$\mu e^{-1} = \mu \left(\frac{1}{2} e^{-(1+\epsilon)k} + \frac{1}{2} e^{-(1-\epsilon)k} \right),$$

and it turns out that k is strictly greater than one. Thus, the selling volume (and thus revenue) increases by setting $\epsilon > 0$. For example when $\epsilon = 0.2$, the selling volume and revenue increases by 2.4%. \square

We highlight that even small deviations from uniform allocations can be profitable, while being quite difficult for consumers or regulators to discover. Notice that under the strategy in Example 2, when the perturbation is randomized for each customer, the total number of each item allocated among all the customers in the market is balanced. Thus a regulator examining aggregate allocation data would not be able to detect the existence of such strategies. Example 2 demonstrates the need for regulators to focus on not only enforcing that sellers follow the stated allocation probabilities, but also to ensure that the seller follows these rules precisely for *each customer*. Thus, effective regulation may need to require that sellers disclose granular, customer-by-customer and transaction-level data so that the regulators may conduct statistical tests to detect unfair strategies.

4.4. Salvage System

In previous sections, we assumed that customers extract zero value from duplicate items received from traditional boxes, and that a customer could not resell items back to the seller. In practice however, some loot box marketplaces are equipped with *salvage* systems, mechanisms by which a customer can trade in unwanted items for currency. Salvage systems are a ubiquitous method for managing customer satisfaction under various sales policies, offering customers a form of recourse against unlikely or unfortunate outcomes (see Fig. EC.3 for an example). In this section, we shall consider loot box selling strategies that allow customers to trade-in or salvage items for a value c . For simplicity, we restrict our attention to the case where the loot box allocates a single item at a time and there is only a single class of items.

The main focus in this setting is to understand the two competing effects that salvage systems have on loot box revenue. On the one hand, the presence of a salvage cost c increases the minimum valuation of any item to at least c , increasing the expected valuation of an item (from $\mathbb{E}[V_i]$ to $\mathbb{E}[\max\{V_i, c\}]$) and thus inducing more purchases. On the other hand, salvage systems return currency to the customer which dilutes the revenue garnered from customer purchases. The results in this section characterize and extend the revenue guarantees of Theorems 2 and 3 to the case when items can be salvaged for some value c . Throughout this section we use the notation $\mathcal{R}(c)$ to denote the optimal revenue of

a strategy with fixed salvage cost c . Note that in the presence of a salvage system, the allocation mechanism for unique box strategies is no longer well-specified. For our results we assume that customers facing a unique loot box strategy are never allocated an item they had previously salvaged, which is the case in the example described in Fig. EC.3.

We first show in Theorem 6 that the introduction of a salvage system by the seller makes both loot box strategies more attractive than separate selling. Specifically, by treating the salvage cost c as a parameter of a loot box strategy, the revenues of both the optimal unique box or traditional box strategies are guaranteed to dominate the revenue of separate selling.

THEOREM 6 (Loot Boxes with Salvage Outperform Separate Selling). *For any N , both the unique box and traditional box strategies with a salvage system dominate separate selling, i.e.,*

$$\max_c \mathcal{R}_{UB}(c) \geq \mathcal{R}_{SS} \text{ and } \max_c \mathcal{R}_{TB}(c) \geq \mathcal{R}_{SS}.$$

Proof of Theorem 6. Let p^* be the optimal price used by separate selling. Now consider a loot box strategy (either unique or traditional) with salvage cost p^* and price p^* . The customer will purchase loot boxes, keeping all the items which they value at p^* or greater and returning the unwanted items for a full refund, until they obtain all items which they value above p^* . Thus, such a loot box induces the same revenue as separate selling, which implies that

$$\max_c \mathcal{R}_{UB}(c) \geq \mathcal{R}_{SS} \text{ and } \max_c \mathcal{R}_{TB}(c) \geq \mathcal{R}_{SS}.$$

□

We emphasize that this result is valid for any finite N . It is well known that the grand bundle is not guaranteed to outperform separate selling for finite N , even though grand bundle selling is asymptotically optimal. Thus, Theorem 6 allows us to pin down the precise relationship between loot box strategies and separate selling and further explains the power and popularity of loot boxes in practice.

We next investigate the revenue of salvage systems in the asymptotic regime. Proposition 5 gives the limiting normalized revenue with respect to a fixed salvage cost c .

PROPOSITION 5 (Revenue and Surplus of Loot Boxes with Salvage Costs.).

Let c be the salvage cost, $\eta = \mathbb{E}[V_i | V_i > c]$, and $\bar{F}(c) = \mathbb{P}(V_i > c)$.

a) *The unique box strategy asymptotically earns*

$$\lim_{N \rightarrow \infty} \mathcal{R}_{UB}(c) = \eta \bar{F}(c),$$

and the limiting normalized customer surplus is $cF(c)$.

b) *The traditional loot box strategy asymptotically earns*

$$\lim_{N \rightarrow \infty} \mathcal{R}_{TB}(c) = \bar{F}(c)(\eta - c) \left(\frac{c}{\eta - c} + e^{-\frac{\eta}{\eta - c}} \right),$$

and the limiting normalized customer surplus is $\bar{F}(c) \left(\eta - c - (2\eta - c)e^{-\frac{\eta}{\eta - c}} \right)$.

This result generalizes the insights derived from Theorem 2 and Theorem 3 to the case with salvage costs. First, note that like before the (asymptotic) revenue of unique box strategies dominates the (asymptotic) revenue of traditional box strategies for any valuation distribution F and salvage cost c . To see this, note since $0 < c \leq \eta$, we may substitute c by $c = q\eta$, for some $q \in [0, 1]$. Plugging in this substitution and rearranging yields:

$$\lim_{N \rightarrow \infty} \mathcal{R}_{TB}(c) \leq \bar{F}(c)\eta \max_{q \in (0, 1]} \left(q + (1 - q)e^{-\frac{1}{1-q}} \right) \leq \bar{F}(c)\eta = \lim_{N \rightarrow \infty} \mathcal{R}_{UB}(c), \quad (2)$$

where the final equality comes from noting $q + (1 - q)e^{-\frac{1}{1-q}}$ is monotone increasing and tends to 1 as $q \rightarrow 1$. Thus $\lim_{N \rightarrow \infty} \mathcal{R}_{UB}(c) \geq \lim_{N \rightarrow \infty} \mathcal{R}_{TB}(c)$. Further, the monotonicity in the maximum in Eq. (2) implies when $\frac{c}{\eta}$ is large (close to 1), the gap in expected revenue between unique box strategies and traditional box strategies is small and generally shrinks from a factor of e ($\frac{c}{\eta} = 0$) monotonically down to 1 ($\frac{c}{\eta} = 1$). Thus when salvage costs are large relative to η , the additional value of employing unique box strategies is diminished. Further, by combining Theorems 2 and 6 and Proposition 5, we obtain a complete ordering of the four strategies in presence of salvage cost:

$$\lim_{N \rightarrow \infty} \mathcal{R}_{GB} = \lim_{N \rightarrow \infty} \max_c \mathcal{R}_{UB}(c) \geq \lim_{N \rightarrow \infty} \max_c \mathcal{R}_{TB}(c) \geq \lim_{N \rightarrow \infty} \mathcal{R}_{SS},$$

Compared to the revenue without salvage, it is worth noting that for the unique boxes, $\bar{F}(c)\eta \leq \mu$, with equality achieved only when $c = 0$. Thus, when N is large it is never optimal for a purely revenue-maximizing seller to use salvage systems with unique boxes. For traditional box, the benefit of introducing a salvage system is distribution-dependent. For example, when V_i is a uniform random variable supported on $[0, 2\mu]$, the asymptotic

optimal traditional box revenue with salvage is 0.517μ , which is 40.4% better than 0.368μ , the revenue without salvage.

Finally, salvage systems are primarily used to improve customer outcomes and overall satisfaction with the system. For unique box strategies, the expected normalized customer surplus under a revenue maximizing unique box with salvage cost c is $cF(c)$. Note the expected normalized customer surplus is monotonically increasing in c . Thus the salvage system enables the unique box seller to balance revenue and customer surplus to their desired proportion.

For traditional box strategies, we find in Theorem 7 that the revenue-maximizing seller may only increase the surplus by at most 1.4%, compared to the case without salvage ($c = 0$).

THEOREM 7 (Salvage System Barely Increases Surplus for (TB)). *The limiting normalized customer surplus of the traditional box strategy with any salvage cost c is at most 1.4% more than the customer surplus of the traditional box strategy with no salvage system ($c = 0$).*

Proof. Recall from Proposition 5 that the customer surplus given salvage c is $\bar{F}(c) \left((\eta - c) - (2\eta - c)e^{-\frac{\eta}{\eta-c}} \right)$. By replacing c with $q\eta$, one can see that the normalized customer surplus can be expressed as

$$\bar{F}(c) \left((\eta - c) - (2\eta - c)e^{-\frac{\eta}{\eta-c}} \right) = \bar{F}(c)\eta \left((1 - q) - (2 - q)e^{-\frac{1}{1-q}} \right). \quad (3)$$

As a special case, when $c = 0$ the normalized surplus is $(1 - \frac{2}{e})\mu = 0.264\mu$ as discussed in Theorem 4. When $c > 0$, note that in the right-hand side of Eq. (3), one can easily see that $\bar{F}(c)\eta \leq \mu$, and $\max_q (1 - q) - (2 - q)e^{-\frac{1}{1-q}} = 0.268$ where the maximum is reached when $q = 0.074$. So using any salvage c may increase the customer surplus by at most 1.4%, i.e., $0.268/0.264 - 1$. In fact, the surplus decreases in most cases.

The following example shows that this bound is tight. Suppose valuations are constant and equal to 1 for all items. Then for any $c < \mu$, $\eta = \mu = 1$. Let $c = 0.074\eta = 0.074$. Then $\bar{F}(c)\eta = 1$, and the normalized surplus is 0.268, which is 1.4% better than 0.264, the surplus from traditional box without salvage. \square

Theorem 7 implies that for traditional boxes, the salvage system may serve as a tool to increase revenue, but does not necessarily improve the customers satisfaction.

5. Numerical Experiments

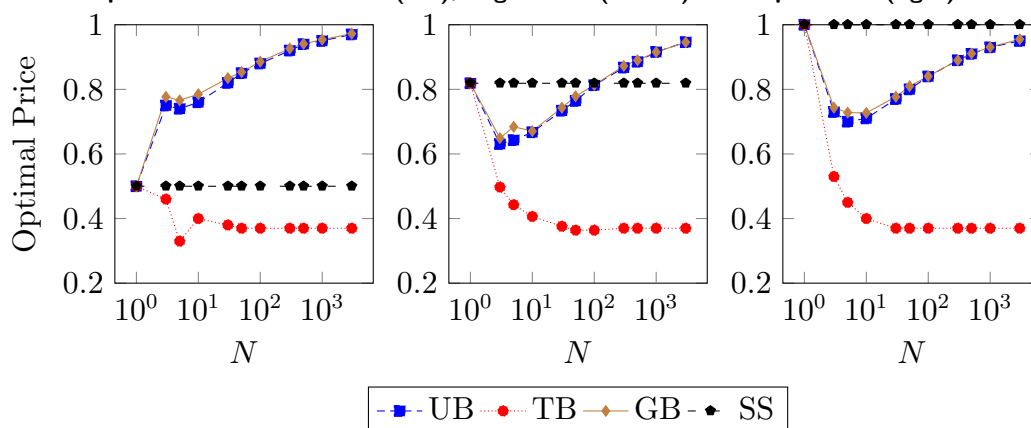
In previous sections, we showed that, asymptotically, unique boxes are optimal and traditional boxes earn a constant fraction of the optimal revenue. In this section, we conduct numerical experiments to demonstrate the efficacy of unique box and traditional box selling for finite N . First, in Section 5.1 we compute and compare the optimal revenues of UB, TB, GB, and SS strategies for typical valuation distributions over a range of values for N that reflect industry practices. We then numerically investigate the impact of budget constraints on the performance of UB, TB, GB and SS strategies in Section 5.2.

In our experiments we consider three possible valuation distributions: the uniform distribution between 0 and 2, the log-normal distribution with log-mean 0 and log-variance 1, and the exponential distribution with mean 1. These distributions are commonly used to model customer valuations, and have been previously studied in Abdallah (2018). For each distribution, we let N range from 1 to 3000, and consider the optimal prices and revenues of the four candidate strategies: UB, TB, GB, and SS. Computation of the revenues is done via simulation by generating 50,000 customer sample paths and using brute force to search for the optimal prices (at 1% accuracy).

5.1. Finite Number of Items

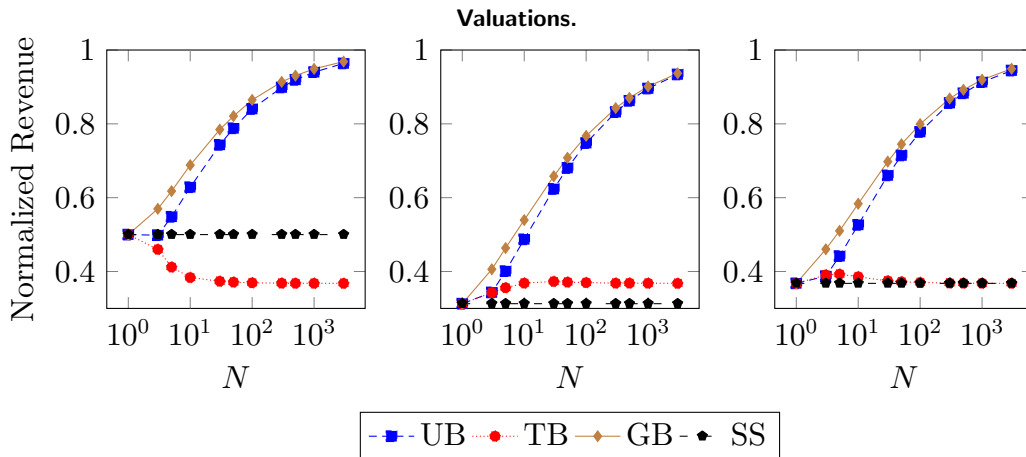
In Fig. 4, we show how the optimal prices of the various strategies change as N increases. Note that for all three distributions that $\mu = 1$ and the optimal price of the traditional box strategy quickly converges to $\frac{1}{e}$ for $N \geq 50$. The optimal price of the unique box strategy also converges to μ for each distribution, albeit at a slower rate.

Figure 4 Optimal Prices for Uniform (left), Log-normal (middle) and Exponential (right) Valuations.



In Fig. 5, we plot the normalized revenues for each strategy. For traditional boxes, the revenue of the optimal policy quickly converge to $\frac{\mu}{e}$ for all three valuation distributions. For unique boxes, the revenue of the optimal policy trends slowly towards it's limit of 1, however it closely follows the revenue of the optimal grand bundle strategy. When $N \geq 100$, the unique box strategy garners more than 70% of the maximum possible revenue and more than 97% of the revenue of the grand bundle. The performance for finite N also depends closely on the distribution. For example, under uniform valuations, the revenue of grand bundle and unique box strategies converge faster than in the case of log-normal valuations. We also note, when N is small, traditional boxes may garner significantly more than $1/e$ under uniform valuations, while under log-normal valuations, it converges to $1/e$ from below.

Figure 5 Normalized Optimal Revenue for Uniform (left), Log-normal (middle) and Exponential (right)



5.2. Budget Constraints

Finally, we look into the impact of budgets on the relative performance of the various strategies. Recall that without budgets, as the catalog of items grows, the normalized revenue of separate selling remains fixed and distribution-dependent whereas grand bundle selling becomes asymptotically optimal (Bakos and Brynjolfsson 1999). However, the fact that customers are budget-constrained may hamper the performance of bundling in practice. In particular, the asymptotic optimality of bundling strategies breaks down when customers vary in their ability to spend on virtual items, i.e., when customers are budget-constrained.

In this subsection, we model the case where customers vary in their budgets to spend on virtual items. Specifically, let each customer's budget be a realization of the random

variable B with distribution F_B . When budgets vary in this fashion, the revenue of grand bundle selling becomes highly dependent on the distribution F_B and may be greatly diminished in some cases, whereas the revenues of separate selling and loot box selling are essentially unchanged.

The (unnormalized) revenue of the optimal grand bundle strategy in the presence of random budgets is:

$$\max_p p \mathbb{P} \left(\sum_{i=1}^N V_i \geq p, B \geq p \right)$$

which is upper bounded by the revenue of a single price strategy in the space of budgets, $\max_p p \mathbb{P}(B \geq p) = \max_p p (1 - F_B(p))$. Note that, like the normalized revenue of separate selling, $\max_p p (1 - F_B(p))$ can be anywhere between 0% to 100% of the expected budget size $\mathbb{E}[B]$ and depends explicitly on F_B .

On the other hand, separate selling and both loot box strategies are quite robust to random fluctuations in customers' budgets. To see this, first consider a customer with budget b facing separate selling. For any fixed price p such a customer simply purchases items until they either obtain all the items for which their valuation exceeds the price, or until they exhaust their budget. In the case where the customer exhausts their budget b , they will purchase $\lfloor \frac{b}{p} \rfloor$ items, spending all of their budget except for possibly some leftover amount less than the price of one item. Thus, for each fixed price p , the revenue of separate selling garnered from each customer is either identical to the case without budgets or close to the budget up to the price of a single item.

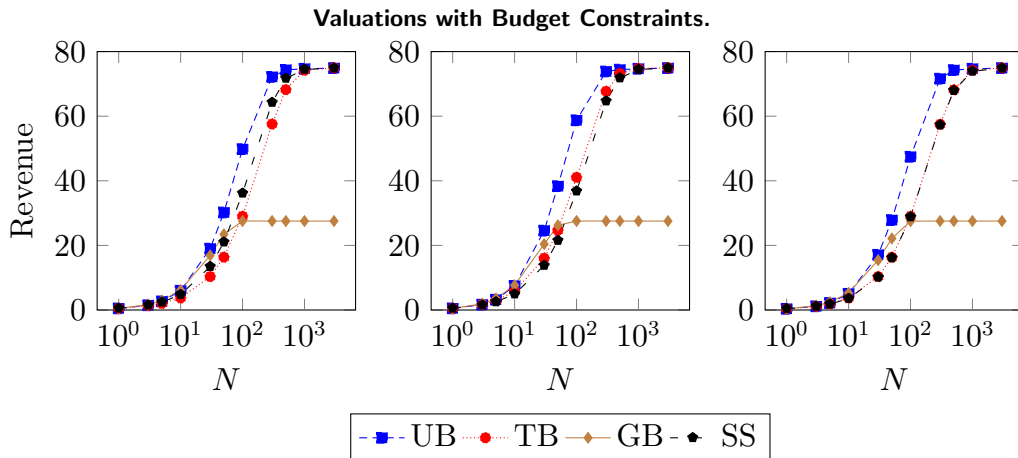
Loot boxes, owing to the sequential nature in which they are purchased, share this property with separate selling. For each customer facing a loot box strategy and fixed price p , they either purchase the same number of boxes as they would in the case without budgets or they spend all of their budget except for possibly some leftover amount less than the price of one box. Combining this observation with Theorem 2 shows that unique boxes *maintain the asymptotic optimality properties of grand bundle selling while remaining effective in the presence of heterogeneous customer budgets*. Thus budget considerations give strong justification for the efficacy and popularity of loot boxes in practice.

In the remainder of this section we study the impact of budgets numerically. First, we specify the distribution of customer budgets. In 2016, 51% of paying customers spent less than \$50 on in-game purchases, and 70% of paying customers spent less than \$100 on in-game purchases (Kunst 2017). To fit this data, we assume that the budget B of a random

customer follows an exponential distribution with a mean of \$75. Under this assumption, 49% of the customers spend less than \$50, and 74% of the customers spend less than \$100.

In Fig. 6, we plot the optimal revenues of each strategy with respect to N , facing customers with exponentially distributed budgets. We use total revenue as the performance measure because the expected revenue now is upper bounded by the expected budget size, \$75. When $N \leq 10$, we note that GB maintains its revenue dominance over the other three strategies. However, as N increases, UB, TB and SS extract almost all the budget, whereas the grand bundle strategy falters, only garnering about 37% the customers budget. This experiment further shows the robustness of loot box strategies in the presence of limited budgets.

Figure 6 (Unnormalized) Optimal Revenue for Uniform (left), Log-normal (middle) and Exponential (right)



6. Conclusions

Our work implies a host of managerial insights for sellers, customers, and regulators of loot boxes. For sellers, we give a thorough analysis of the profitability of loot boxes, yielding guidelines for how to design and price loot boxes so as to maximize revenue. We show that the unique box strategy is asymptotically optimal, whereas the traditional box can garner only 36.7% of the maximum revenue. These results hold in the cases where loot boxes allocate multiple items as well as when the items are heterogeneous. When customers are budget-constrained the implication is even clearer: unique box selling stands above the other three strategies, retaining the asymptotic efficiency properties of grand bundle selling while remaining robust to fluctuations in customers ability to spend. Finally, and

perhaps surprisingly, even in the case where items come from multiple classes, designing the allocation rule for either loot box strategy is easy: a simple uniform allocation policy is effective and asymptotically optimal.

From the customer's perspective, we show that the traditional loot box strategy is preferred and the unique box strategy does not yield any customer surplus. Further, we show that the introduction of a salvage system has surprisingly little affect on customer surplus when facing traditional loot boxes, with a potential gain of only 1.4%. We also show that customers may be at risk to seller manipulation. Specifically, if the seller deviates from the announced allocation probabilities, then they are capable of making more revenue, even when the allocation probabilities are correct in expectation. Thus, it is essential for regulators to protect consumers against such a scenario. In fact, we show that the regulator must check each customer's allocations individually to properly ensure that the seller is being truthful. These facts together show that ex post analysis of allocations may be insufficient in detecting fraudulent behavior on behalf of the seller and suggest that regulatory bodies must inspect customer-dependent data streams or loot box mechanisms at their source implementation to ensure truthful behavior.

Finally, while loot box selling has recently gained attention in the domains of psychology and policy-making, there is a distinct lack of academic work which analyzes loot boxes from a revenue management perspective. Our work breaks ground on this topic, but leaves open several interesting avenues for future work. One particularly fruitful direction for future work may consider loot boxes under richer customer valuation models. In our work, valuations for the items are assumed to be a priori identical for each customer, and independent of the customer's current collection. Follow-up work may consider some form of structured dependency among the customer valuations for the items, e.g., they are modulated by a common customer-specific random factor. One may also consider the case where valuations are no longer additive, but submodular or supermodular in the items currently owned. For example, customers may highly value the very last unowned item since it is necessary to complete their collection. In this case, the revenue from a traditional box strategy may increase dramatically, since it takes many purchases in expectation to collect all the items. On the other hand, if many of the items are substitutable, the customers might be less interested to open the last few boxes since their additional value is marginal relative to the current collection.

Further, our model does not specify the time scale on which the items are acquired. Customer behavior may be more complicated when the utility is proportional to how often the item is used. Within this setting, it is possible that customers have higher valuations for items acquired earlier since the player can use those items for longer. One may also consider loot box selling when the catalog of items is increased dynamically over time. For such models, the cost of introducing new items may be non-trivial and increase with the catalog size as the complexity required to maintain game balance grows. The sequence in which the new items are added to the catalog may also have a non-trivial impact on the revenue garnered. Finally, in connection with the ongoing debate in the media and governments, it would be interesting to consider loot box pricing and design problems under various legal or fairness considerations.

Acknowledgments

We greatly appreciate the feedback from the three anonymous reviews, the AE, and the DE, all of whom helped improve the paper significantly.

References

- Abdallah, Tarek. 2018. On the benefit (or cost) of large-scale bundling. *Production and Operations Management* **0**(0).
- Abdallah, Tarek, Arash Asadpour, Josh Reed. 2019. Large-scale bundle size pricing: A theoretical analysis. *Working Paper, Northwestern University* .
- Apple(March-8). 2018. App store review guidelines URL <https://developer.apple.com/app-store/review/guidelines/>.
- Babaioff, Moshe, Nicole Immorlica, Brendan Lucier, S. Weinberg. 2020. A simple and approximately optimal mechanism for an additive buyer. *Journal of the ACM* **67** 1–40. doi:10.1145/3398745.
- Bakos, Yannis, Erik Brynjolfsson. 1999. Bundling information goods: Pricing, profits, and efficiency. *Management Science* **45**(12) 1613–1630. doi:10.1287/mnsc.45.12.1613.
- Batchelor, James. 2017. NPD: Loot box controversy having no impact on game sales URL <https://www.gamesindustry.biz/articles/2017-10-25-npd-loot-box-controversy-having-no-impact-on-sales-potential>.
- Briest, Patrick, Shuchi Chawla, Robert Kleinberg, S. Weinberg. 2014. Pricing lotteries. *Journal of Economic Theory* **156**. doi:10.1016/j.jet.2014.04.011.
- Briest, Patrick, Heiko Röglin. 2010. The power of uncertainty: Bundle-pricing for unit-demand customers. 47–58. doi:10.1007/978-3-642-18318-8_5.

- Chen, R., E. Gal-Or, P. Roma. 2014. Opaque distribution channels for competing service providers: Posted price vs. name-your-own-price mechanisms. *Operations research* **62**(4) 733–750.
- Chen, Zhengxing, Su Xue, John Kolen, Navid Aghdaie, Kazi A Zaman, Yizhou Sun, Magy Seif El-Nasr. 2017. Eomm: An engagement optimized matchmaking framework. *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1143–1150.
- Drummond, Aaron, James D. Sauer. 2018. Video game loot boxes are psychologically akin to gambling. *Nature Human Behaviour* **2**(8) 530–532. doi:10.1038/s41562-018-0360-1. URL <https://doi.org/10.1038/s41562-018-0360-1>.
- Elmachtoub, Adam N., Michael L. Hamilton. 2020. The power of opaque products in pricing. *Working Paper, Columbia University* .
- Elmachtoub, Adam N, Yehua Wei, Yeqing Zhou. 2020. Retailing with opaque products. *Working Paper, Columbia University* .
- Elmachtoub, Adam N, David Yao, Yeqing Zhou. 2019. The value of flexibility from opaque selling. *Working Paper, Columbia University* .
- Fay, S., J. Xie. 2008. Probabilistic goods: A creative way of selling products and services. *Marketing Science*, vol. 27. 674–690.
- Ferreira, Kris, Joel Goh. 2020. Assortment rotation and the value of concealment. *Management Science* doi:10.1287/mnsc.2020.3610.
- Fingas, Jon. 2018. South Korea fines game studios over deceptive loot box odds. <https://www.engadget.com/2018/04/10/south-korea-fines-game-studios-over-loot-boxes>.
- Gallego, G., G. Iyengar, R. Phillips, A. Dubey. 2004. Managing flexible products on a network .
- Gallego, Guillermo, Anran li, Van-Anh Truong, Xinshang Wang. 2020. Approximation algorithms for product framing and pricing. *Operations Research* **68**. doi:10.1287/opre.2019.1875.
- Gallego, Guillermo, Robert Phillips. 2004. Revenue management of flexible products. *Manufacturing & Service Operations Management* **6** 321–337. doi:10.1287/msom.1040.0054.
- Golrezaei, Negin, Vahideh Manshadi, Vahab Mirrokni. 2018. Two-stage pandora’s box for product ranking. *Working Paper, Massachusetts Institute of Technology* .
- Google(July-23). 2019. Monetization and ads URL <https://play.google.com/about/monetization-ads/>.
- Gut, Allan. 2009. *Stopped random walks*. Springer.
- Hart, Sergiu, Noam Nisan, et al. 2014. *How good are simple mechanisms for selling multiple goods?*. Citeseer.
- Hart, Sergiu, Philip J. Reny. 2015. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics* **10**(3) 893–922.

- Holt, Kris. 2019. The FTC will hold a public workshop on loot boxes in August. <https://www.engadget.com/2019/04/08/ftc-loot-boxes-investigation-public-workshop/>.
- Huang, Xiao, Greys Sošić, Gregory Kersten. 2017. Selling through priceline? on the impact of name-your-own-price in competitive market. *IISE Transactions* **49**(3) 304–319.
- Huang, Yan, Stefanus Jasin, Puneet Manchanda. 2018. 'level up': Leveraging skill and engagement to maximize player game-play in online video games. Available at SSRN: <https://ssrn.com/abstract=2885082> .
- Jerath, Kinshuk, Serguei Netessine, Senthil K Veeraraghavan. 2010. Revenue management with strategic customers: Last-minute selling and opaque selling. *Management Science* **56**(3) 430–448.
- Jiang, Yabing. 2007. Price discrimination with opaque products. *Journal of Revenue and Pricing Management* **6** 118–134. doi:10.1057/palgrave.rpm.5160073.
- Jiao, Yifan, Christopher S Tang, Jingqi Wang. 2020. Opaque selling in player-versus-player games. Available at SSRN [3558774](https://ssrn.com/abstract=3558774) .
- Kunst, Alexander(January-19). 2017. Amount spent on in-game purchases according to gamers in the united states as of december 2016 URL <https://www.statista.com/statistics/662618/amount-spent-in-game-purchases-us-gamers/>.
- Lorden, Gary. 1970. On excess over the boundary. *The Annals of Mathematical Statistics* 520–527.
- Ma, Will, David Simchi-Levi. 2016. Reaping the benefits of bundling under high production costs. *Working Paper, Columbia University* .
- Perez, Matt. 2018. Report: For the first time, mobile will make up half of all gaming revenue URL <https://www.forbes.com/sites/mattperez/2018/04/30/report-for-the-first-time-mobile-will-make-up-half-of-all-gaming-revenue/#50f486a1586a>.
- Romm, Tony, Craig Timberg. 2019. Video game 'loot boxes' would be outlawed in many games under forthcoming federal bill URL https://www.washingtonpost.com/technology/2019/05/08/video-game-loot-boxes-would-be-outlawed-many-games-under-forthcoming-federal-bill/?noredirect=on&utm_term=.ffc0b8ac8b73.
- Shapiro, Dmitry, Xianwen Shi. 2008. Market segmentation: The role of opaque travel agencies. *Journal of Economics & Management Strategy* **17**(4) 803–837.
- Sheng, Lifei, Mahesh Nagarajan, Christopher Ryan, Yuan Cheng, Chunyang Tong. 2020. Incentivized actions in freemium games. *Working Paper, University of Houston - Clear Lake* .
- Sinclair, Brendan. 2019. Loot boxes a matter of "life or death," says researcher URL <https://www.gamesindustry.biz/articles/2019-08-08-loot-boxes-a-matter-of-life-or-death>.

- Tassi, Paul. 2017. China Has Forced Blizzard To Reveal Exact 'Overwatch' And 'Hearthstone' Drop Rates. <https://www.forbes.com/sites/insertcoin/2017/05/05/china-has-forced-blizzard-to-reveal-exact-overwatch-and-hearthstone-drop-rates/#615b69035a14>.
- Tassi, Paul. 2018. EA now seems legitimately terrified of loot boxes after 'battlefront 2' URL <https://www.forbes.com/sites/insertcoin/2018/04/17/ea-now-seems-legitimately-terrified-of-loot-boxes-after-battlefront-2/#2d7215b7648e>.
- Wang, Ruxian, Ozge Sahin. 2017. The impact of consumer search cost on assortment planning and pricing. *Management Science* **64**(8) 3649–3666.
- Webb, Kevin. 2017. Regulators from more than a dozen countries are looking to crack down on 'loot boxes,' a controversial video gaming practice that could be too much like gambling. URL <https://www.businessinsider.com/loot-boxes-european-regulation-2018-9>.
- Wright, Mike. 2018. Video gamers will be spending \$50 billion on 'gambling-like' loot box features by 2022, according to analysts URL <https://www.telegraph.co.uk/technology/2018/04/17/video-gamers-will-spending-50-billion-gambling-like-loot-box/>.
- Xiao, Yongbo, Jian Chen. 2014. Evaluating the potential effects from probabilistic selling of similar products. *Naval Research Logistics (NRL)* **61**(8) 604–620.
- Zendle, David, Paul Cairns. 2018. Video game loot boxes are linked to problem gambling: Results of a large-scale survey. *PloS one* **13**(11) e0206767.

Online Appendix: Loot Box Pricing and Design

Appendix A: Omitted Proofs

A.1. Omitted Proofs from Section 2

Proof of Theorem 1.

(a) *Unique Box.* Without loss of generality, we index the items by the order in which they are allocated, i.e., the i^{th} unique box yields the i^{th} item which is valued at V_i . Let $x_k = \sum_{i=1}^k (v_i - p)$, which is the net utility of buying k unique boxes. Note that the myopic customer will buy the first unique box if $U_1 = x_N/N \geq 0$. The customer will continue to purchase until the first time $\tau + 1$ where $U_{\tau+1} < 0$, i.e., until

$$U_{\tau+1} = \frac{\sum_{i \in [N] \setminus S_{\tau+1}} v_i}{N - |S_{\tau+1}|} - p = \frac{\sum_{i=\tau+1}^N v_i}{N - \tau} - p = \frac{x_N - x_\tau}{N - \tau} < 0. \quad (\text{EC.1})$$

When this condition is met, the customer stops with τ items in total and has total utility of x_τ . For clarity, we set $U_{N+1} < 0$ to include the case that $\tau = N$. From Eq. EC.1, we find that τ is also the first time that $x_\tau > x_N$ (or $\tau = N$ if $x_t \leq x_N$ for all t) since

$$x_\tau = \sum_{i=1}^{\tau} (v_i - p) = x_N - \sum_{i=\tau+1}^N (v_i - p) = x_N - (N - \tau)U_{\tau+1} \geq x_N. \quad (\text{EC.2})$$

Hence, the net utility of the myopic strategy, x_τ , is at least x_N .

Observe that an upper bound on the maximum possible utility of any purchasing strategy is the utility of the clairvoyant strategy, denoted by M_N , that stops when customer utility is maximized, i.e.,

$$M_N = \max_{k \in 0 \cup [N]} \sum_{i=1}^k (v_i - p),$$

where $\sum_{i=1}^0 (v_i - p) = 0$. For a *random* customer, M_N is equivalent to the maximum value of a random walk on $\{X_k\}_{k=0}^N$, where $X_k = \sum_{i=1}^k (V_i - p)$. By Theorem 2.12.1 in Gut (2009), $\lim_{N \rightarrow \infty} M_N/N$ converges to $\max(0, \mu - p)$ almost surely. Further by the strong law of large numbers, $\lim_{N \rightarrow \infty} X_N/N$ converges to $\mu - p$ almost surely, which implies that the expected normalized net utility of a myopic customer converges to $\max(0, \mu - p)$ almost surely, because a myopic customer always garners at least $\max(0, X_N)$ utility. Therefore, the expected normalized net utility of the myopic strategy and optimal strategies must also converge to $\max(0, \mu - p)$ almost surely.

(b) *Traditional Box.* Let $\tau + 1$ be the first period where a myopic customer decides not purchase an item, meaning the customer purchased exactly τ traditional boxes. This means that

$$U_{\tau+1} = \frac{\sum_{i \in [N] \setminus S_{\tau+1}} v_i}{N} - p < 0. \quad (\text{EC.3})$$

However, it is clear that the utility U_t is non-increasing in t . If a duplicate is received in period t , then $U_t = U_{t+1}$. If a new item is received, then the customer now values that item as 0 in the future and their expected valuation of a traditional loot box decreases. Therefore, once $\tau + 1$ is reached, a customer will never see a traditional box that offers a positive utility even if they continue purchasing indefinitely. Thus, a myopic strategy is optimal. \square

A.2. Omitted Proofs from Section 3

Proof of Theorem 2. The proof works by constructing a sequence of prices, p_N , and showing that $\mathcal{R}_{UB}(p_N)$ is greater than a term that converges to μ as N goes to ∞ . Since μ is the maximum possible normalized revenue by Proposition 1, this then implies that $\lim_{N \rightarrow \infty} \mathcal{R}_{UB}(p_N) = \lim_{N \rightarrow \infty} \mathcal{R}_{UB} = \mu$. We shall rely on the random walk from the proof of Theorem 1(a), $\{X_j\}_{j=0}^{\infty}$ where $X_j := \sum_{i=1}^j (V_i - p_N)$ and $X_0 = 0$. Without loss of generality, we assume the items are indexed so that the i^{th} item a customer receives from the i^{th} unique box is valued at V_i . Let the random variable τ_N denote the number of boxes purchase, and recall from Eq. (EC.2) that τ_N is also the first time that $X_{\tau_N} > X_N$, or $\tau_N = N$ if $X_t \leq X_N$ for all t . Note that $\mathcal{R}_{UB}(p_N) = p_N \mathbb{E}[\tau_N]/N$. Also note that since X_N is not known to the seller, τ_N is not a stopping time. (However, it is a stopping time from the perspective of the customer.) We shall show that for a sequence of prices that tend to μ , that τ_N tends to N which implies our result.

For some $\mu > \epsilon_N > 0$ to be optimized later, let $p_N = \mu - \epsilon_N$. We shall compare τ_N to an actual stopping time $\bar{\tau}_N$, which is the first time $\{X_i\}$ crosses the threshold $(1 - k_N)N\epsilon_N$, where $1 > k_N > 0$ shall be optimized later. Note that if we condition on the event that $X_N \geq (1 - k_N)N\epsilon_N$, then we know that $\tau_N \geq \bar{\tau}_N$ since $\{X_i\}$ must hit $(1 - k_N)N\epsilon_N$ before hitting X_N . Therefore,

$$\begin{aligned} \mathbb{E}[\tau_N] &\geq \mathbb{E}[\tau_N \mathbb{I}_{X_N \geq (1-k_N)N\epsilon_N}] \\ &\geq \mathbb{E}[\bar{\tau}_N \mathbb{I}_{X_N \geq (1-k_N)N\epsilon_N}] \\ &= \mathbb{E}[\bar{\tau}_N] - \mathbb{E}[\bar{\tau}_N \mathbb{I}_{X_N \in [0, (1-k_N)N\epsilon_N)}] - \mathbb{E}[\bar{\tau}_N \mathbb{I}_{X_N < 0}]. \end{aligned} \quad (\text{EC.4})$$

We proceed by lower bounding Eq. (EC.4) term by term, beginning with $\mathbb{E}[\bar{\tau}_N]$. Since $\bar{\tau}_N$ is a stopping time, by Wald's equation we know that

$$\mathbb{E}[X_{\bar{\tau}_N}] = \mathbb{E}\left[\sum_{i=1}^{\bar{\tau}_N} (V_i - p_N)\right] = \mathbb{E}[\bar{\tau}_N] \mathbb{E}[V_i - p_N] = \mathbb{E}[\bar{\tau}_N] \epsilon_N. \quad (\text{EC.5})$$

Rearranging (EC.5), we have

$$\mathbb{E}[\bar{\tau}_N] = \frac{\mathbb{E}[X_{\bar{\tau}_N}]}{\epsilon_N} \geq \frac{(1 - k_N)N\epsilon_N}{\epsilon_N} = (1 - k_N)N, \quad (\text{EC.6})$$

where the inequality follows from the definition of $\bar{\tau}_N$.

Next, we provide an upper bound for the second term in (EC.4), $\mathbb{E}[\bar{\tau}_N \mathbb{I}_{X_N \in [0, (1-k_N)N\epsilon_N)}]$. This term corresponds to the case that $X_N \in [0, (1 - k_N)N\epsilon_N)$. To derive an upper bound, we suppose that $\{X_i\}$ has not hit $(1 - k_N)N\epsilon_N$ after N steps, and further assume the worst case that $X_N = 0$. In this case, it is a fresh random walk starting from 0. We first note that for a discrete random walk crossing a threshold, by (Lorden 1970) Theorem 1 we have

$$\mathbb{E}[X_{\bar{\tau}_N} - (1 - k_N)N\epsilon_N] \leq \frac{\mathbb{E}[\max(V_i - \mu + \epsilon_N, 0)^2]}{\epsilon_N} \leq \frac{\mathbb{E}[(V_i - \mu + \epsilon_N)^2]}{\epsilon_N} = \frac{\mathbb{E}[V_i - \mu + \epsilon_N]^2 + \sigma^2}{\epsilon_N} \leq \frac{\mu^2 + \sigma^2}{\epsilon_N}$$

which implies that

$$\mathbb{E}[X_{\bar{\tau}_N}] \leq (1 - k_N)N\epsilon_N + \frac{\mu^2 + \sigma^2}{\epsilon_N}.$$

Hence, by Wald's equation (see (EC.5)), in expectation it takes at most another $(1 - k_N)N + \frac{\mu^2 + \sigma^2}{\epsilon_N^2}$ steps to hit $(1 - k_N)N\epsilon_N$ if $X_N = 0$. Thus,

$$\mathbb{E}[\bar{\tau}_N | X_N \in [0, (1 - k_N)N\epsilon_N]] \leq N + (1 - k_N)N + \frac{\mu^2 + \sigma^2}{\epsilon_N^2}. \quad (\text{EC.7})$$

The probability that $X_N \in [0, (1 - k_N)N\epsilon_N]$ can be upper bounded using Chebyshev's Inequality,

$$\mathbb{P}(X_N \in [0, (1 - k_N)N\epsilon_N]) \leq \mathbb{P}(X_N < (1 - k_N)N\epsilon_N) \leq \frac{\sigma^2}{k_N^2 \epsilon_N^2 N}. \quad (\text{EC.8})$$

Combining (EC.7) and (EC.8), we have

$$\mathbb{E}[\bar{\tau}_N \mathbb{I}_{X_N \in [0, (1 - k_N)N\epsilon_N]}] \leq \left(N + (1 - k_N)N + \frac{\mu^2 + \sigma^2}{\epsilon_N^2} \right) \frac{\sigma^2}{k^2 \epsilon_N^2 N} = \frac{(2 - k_N)\sigma^2}{k_N^2 \epsilon_N^2} + \frac{(\mu^2 + \sigma^2)\sigma^2}{k_N^2 \epsilon_N^4 N}. \quad (\text{EC.9})$$

Next, we provide an upper bound for the third term in (EC.4), $\mathbb{E}[\bar{\tau}_N \mathbb{I}_{X_N < 0}]$. This term corresponds to the case that $X_N < 0$. To derive an upper bound, we suppose that $\{X_i\}$ has not hit $(1 - k_N)N\epsilon_N$ after N steps, and further assume the worst case that $X_N = -Np_N$ (since $V_i \geq 0$). Following the same logic as (EC.7),

$$\mathbb{E}[\bar{\tau}_N | X_N < 0] \leq N + \frac{Np + (1 - k_N)N\epsilon_N}{\epsilon_N} + \frac{\mu^2 + \sigma^2}{\epsilon_N^2}. \quad (\text{EC.10})$$

As before, the probability that $X_N < 0$ can also be upper bounded using Chebyshev's Inequality,

$$\mathbb{P}(X_N < 0) \leq \frac{\sigma^2}{\epsilon_N^2 N}. \quad (\text{EC.11})$$

Combining (EC.10) and (EC.11) yields

$$\mathbb{E}[\bar{\tau}_N \mathbb{I}_{X_N < 0}] \leq \left(N + \frac{Np + (1 - k_N)N\epsilon_N}{\epsilon_N} + \frac{\mu^2 + \sigma^2}{\epsilon_N^2} \right) \frac{\sigma^2}{\epsilon_N^2 N} = \left(\frac{\mu}{\epsilon_N} + 1 - k_N \right) \frac{\sigma^2}{\epsilon_N^2} + \frac{(\mu^2 + \sigma^2)\sigma^2}{\epsilon_N^4 N}. \quad (\text{EC.12})$$

Plugging Eqs. (EC.6), (EC.9) and (EC.12) into the right hand side of Eq. (EC.4) yields

$$\mathbb{E}[\tau_N] \geq N \left(1 - k_N - \frac{2\sigma^2}{k_N^2 \epsilon_N^2 N} + \frac{\sigma^2}{k_N \epsilon_N^2 N} - \frac{\sigma^2 \mu}{\epsilon_N^3 N} - \frac{\sigma^2}{\epsilon_N^2 N} + \frac{\sigma^2 k_N}{\epsilon_N^2 N} - \frac{(\mu^2 + \sigma^2)\sigma^2}{k_N^2 \epsilon_N^4 N^2} - \frac{(\mu^2 + \sigma^2)\sigma^2}{\epsilon_N^4 N^2} \right).$$

Now we can lower bound the normalized revenue of a unique box strategy,

$$\begin{aligned} \mathcal{R}_{UB} &\geq (\mu - \epsilon_N) \frac{\mathbb{E}[\tau_N]}{N} \\ &\geq \mu \left(1 - \frac{\epsilon_N}{\mu} \right) \left(1 - k_N - \frac{2\sigma^2}{k_N^2 \epsilon_N^2 N} + \frac{\sigma^2}{k_N \epsilon_N^2 N} - \frac{\sigma^2 \mu}{\epsilon_N^3 N} - \frac{\sigma^2}{\epsilon_N^2 N} + \frac{\sigma^2 k_N}{\epsilon_N^2 N} - \frac{(\mu^2 + \sigma^2)\sigma^2}{k_N^2 \epsilon_N^4 N^2} - \frac{(\mu^2 + \sigma^2)\sigma^2}{\epsilon_N^4 N^2} \right). \end{aligned}$$

Choosing $\epsilon_N = \mu N^{-1/5}$, and $k_N = N^{-1/5}$, we have

$$\mathcal{R}_{UB} \geq \mu (1 - N^{-1/5}) \left(1 - (1 + \frac{2\sigma^2}{\mu^2}) N^{-1/5} - \frac{\sigma^2}{\mu^2} N^{-3/5} - \frac{\sigma^4}{\mu^4} N^{-4/5} - (\frac{\sigma^2}{\mu^2} + \frac{\sigma^4}{\mu^4}) N^{-6/5} \right).$$

Taking the limit of both sides gives

$$\lim_{N \rightarrow \infty} \mathcal{R}_{UB} \geq \mu.$$

Combined with the fact that $\mathcal{R}_{UB} \leq \mu$ from Proposition 1, we conclude that $\lim_{N \rightarrow \infty} \mathcal{R}_{UB} = \mu$. \square

Proof of Theorem 3. Consider a random walk for N steps, $\{Y_j\}_{j=0}^N$, where $Y_j = \sum_{i=j+1}^N V_i$ for $j = 0, \dots, N-1$ and $Y_N = 0$. Without loss of generality, we assume the items are indexed so that the i^{th} unique item a customer receives from purchasing traditional loot boxes is valued at V_i . Therefore, every time the customer receives the i^{th} unique item, their valuation for the traditional box becomes Y_i/N .

Similar to our proof of Theorem 2, we construct a sequence of prices p_N such that $\lim_N p_N \rightarrow \frac{\mu}{e}$ and show the expected number of traditional loot boxes purchased by a customer at price p_N tends to N . Let the random variable τ_N denote the number of unique items acquired, and recall from Eq. (EC.3) that τ_N is also the first time $Y_{\tau_N}/N - p_N < 0$. Note that τ_N is well defined since $Y_N = 0$.

The number of traditional loot boxes a customer must have purchased to acquire τ_N unique items is the sum of τ_N independent geometric random variables, $\text{Geo}(1) + \text{Geo}(\frac{N-1}{N}) + \dots + \text{Geo}(\frac{N-\tau_N+1}{N})$. The revenue under price p_N is then,

$$\begin{aligned} \mathcal{R}_{TB}(p_N) &= \frac{1}{N} \mathbb{E} \left[p_N \left(\text{Geo}(1) + \text{Geo} \left(\frac{N-1}{N} \right) + \dots + \text{Geo} \left(\frac{N-\tau_N+1}{N} \right) \right) \mathbb{I}_{\tau_N \geq 1} \right] \\ &= p_N \mathbb{E} \left[\frac{1}{N} \left(1 + \frac{N}{N-1} + \dots + \frac{N}{N-\tau_N+1} \right) \mathbb{I}_{\tau_N \geq 1} \right] \\ &= p_N \mathbb{E} \left[(\log(N) + \gamma + \zeta_N - \log(N - \tau_N + 1) - \gamma - \zeta_{N-\tau_N+1}) \mathbb{I}_{\tau_N \geq 1} \right] \\ &= p_N \mathbb{E} \left[\left(-\log \frac{N - \tau_N + 1}{N} + \zeta_N - \zeta_{N-\tau_N+1} \right) \mathbb{I}_{\tau_N \geq 1} \right], \end{aligned} \quad (\text{EC.13})$$

where the third equality follows from the well known expression for the harmonic numbers, $\sum_{i=1}^k \frac{1}{i} = \log k + \gamma + \zeta_k$, with $\{\zeta_k\}$ converges to 0 from above, and γ is the Euler-Mascheroni constant.

First we bound $\mathbb{E}[N - \tau_N + 1]$. Let us define the monotonically increasing random walk $\{C_j\}_{j=0}^\infty$ such that (i) $\{C_j\}_{j=0}^N = \{Y_{N-j}\}_{j=0}^N$, i.e., $C_0 = 0$, $C_1 = V_N$, $C_2 = V_N + V_{N-1}$, \dots , $C_N = V_N + \dots + V_1$ and (ii) $C_j = C_N + \sum_{k=N-j+1}^0 V_k$ for $j = N+1, N+2, \dots$ where $V_0, V_{-1}, V_{-2}, \dots$ are virtual random variables that are i.i.d. samples from F . Let r_N be the first time that $\{C_j\}_{j=1}^\infty$ is at least Np_N . By definition of τ_N , note that when $\tau_N \geq 1$, $r_N = N - \tau_N + 1$. Since r_N is the first passage time when $C_j \geq Np_N$, it follows by the well known inspection paradox that $\mathbb{E}[C_{r_N} - C_{r_N-1}] = \frac{\mathbb{E}[V_i^2]}{\mathbb{E}[V_i]} = \frac{\mu^2 + \sigma^2}{\mu}$. Using this fact together with Wald's equation, $\mathbb{E}[C_{r_N}] = \mathbb{E}[r_N]\mu$, we have

$$\mathbb{E}[r_N] = \frac{\mathbb{E}[C_{r_N}]}{\mu} \in \left[\frac{Np}{\mu}, \frac{Np}{\mu} + 1 + \frac{\sigma^2}{\mu^2} \right]. \quad (\text{EC.14})$$

Now we can construct a lower bound for $\mathcal{R}_{TB}(p_N)$,

$$\mathcal{R}_{TB}(p_N) = p_N \mathbb{E} \left[\left(-\log \frac{N - \tau_N + 1}{N} + \zeta_N - \zeta_{\tau_N} \right) \mathbb{I}_{\tau_N \geq 1} \right] \quad (\text{Eq. (EC.13)})$$

$$\geq p_N \mathbb{E} \left[\left(-\log \frac{N - \tau_N + 1}{N} \right) \mathbb{I}_{\tau_N \geq 1} \right] \quad (\{\zeta_k\} \text{ monotone dec.})$$

$$= p_N \mathbb{E} \left[\left(-\log \frac{r_N}{N} \right) \mathbb{I}_{r_N \leq N} \right] \quad (r_N \leq N \iff \tau_N \geq 1) \quad (\text{EC.15})$$

$$\geq p_N \mathbb{E} \left[\left(-\log \frac{r_N}{N} \right) \right] \quad (\text{EC.16})$$

$$\geq -p_N \log \mathbb{E} \left[\frac{r_N}{N} \right] \quad (\text{Jensen's Inequality})$$

$$\geq -p_N \log \left(\frac{p_N}{\mu} + \frac{1 + \frac{\sigma^2}{\mu^2}}{N} \right) \quad (\text{Eq. (EC.14)}) \quad (\text{EC.17})$$

where Eq. (EC.15) follows from the fact that $-\log \frac{r_N}{N} < 0$ when $r_N > N$. Setting $p_N = \frac{\mu}{e}$ yields

$$\mathcal{R}_{TB}(p_N) \geq \frac{\mu}{e} \log \left(\frac{1}{\frac{1}{e} + \frac{1 + \frac{\sigma^2}{\mu^2}}{N}} \right) \quad (\text{EC.18})$$

which is our desired guarantee.

We now upper bound the revenue, $\mathcal{R}_{TB}(p_N)$. Consider the event that $\frac{r_N}{N} \leq (1 - \epsilon_N) \frac{p_N}{\mu}$ for some small $1 > \epsilon_N > 0$, which is an ingredient of our proof. We can upper bound the probability of such an event by,

$$\begin{aligned}
\mathbb{P}\left(\frac{r_N}{N} \leq (1 - \epsilon_N) \frac{p_N}{\mu}\right) &= \mathbb{P}\left(r_N \leq (1 - \epsilon_N) \frac{Np_N}{\mu}\right) \\
&= \mathbb{P}\left(\sum_{t=N-(1-\epsilon_N)\frac{Np_N}{\mu}}^N V_t \geq Np_N\right) \\
&= \mathbb{P}\left(\frac{\sum_{t=N-(1-\epsilon_N)\frac{Np_N}{\mu}}^N V_t}{(1 - \epsilon_N) \frac{Np_N}{\mu}} \geq \frac{1}{1 - \epsilon_N} \mu\right) \\
&\leq \mathbb{P}\left(\left|\frac{\sum_{t=N-(1-\epsilon_N)\frac{Np_N}{\mu}}^N V_t}{(1 - \epsilon_N) \frac{Np_N}{\mu}} - \mu\right| \geq \frac{\epsilon_N}{1 - \epsilon_N} \mu\right) \\
&\leq \frac{\sigma^2}{(1 - \epsilon_N) \frac{Np_N}{\mu} \frac{\epsilon_N^2}{(1 - \epsilon_N)^2} \mu^2} && \text{(Chebyshev's Inequality)} \\
&= \frac{\sigma^2}{\frac{\epsilon_N^2}{(1 - \epsilon_N)} Np_N \mu}. && \text{(EC.19)}
\end{aligned}$$

Now we can upper bound the revenue from (TB) when using price p_N by

$$\begin{aligned}
\mathcal{R}_{TB}(p_N) &= p_N \mathbb{E}\left[\left(-\log \frac{N - \tau_N + 1}{N} + \zeta_N - \zeta_{\tau_N}\right) \mathbb{I}_{\tau_N \geq 1}\right] \\
&\leq p_N \mathbb{E}\left[\left(-\log \frac{N - \tau_N + 1}{N} + \zeta_N\right) \mathbb{I}_{\tau_N \geq 1}\right] \\
&= p_N \mathbb{E}\left[\left(-\log \frac{r_N}{N} + \zeta_N\right) \mathbb{I}_{r_N \leq N}\right] \\
&\leq p_N \zeta_N + p_N \mathbb{E}\left[\left(-\log \frac{r_N}{N}\right) \mathbb{I}_{r_N \leq N}\right] \\
&= p_N \zeta_N + p_N \mathbb{E}\left[\left(-\log \frac{r_N}{N}\right) \mathbb{I}_{r_N \in [N(1-\epsilon_N)\frac{p_N}{\mu}, N]}\right] + p_N \mathbb{E}\left[\left(-\log \frac{r_N}{N}\right) \mathbb{I}_{r_N \in [1, N(1-\epsilon_N)\frac{p_N}{\mu}]}\right] \\
&\leq p_N \zeta_N + p_N \max\left\{-\log\left((1 - \epsilon_N) \frac{p_N}{\mu}\right), 0\right\} + p_N \mathbb{E}\left[\left(-\log \frac{r_N}{N}\right) \mathbb{I}_{r_N \in [1, N(1-\epsilon_N)\frac{p_N}{\mu}]}\right] \\
&\leq p_N \zeta_N + p_N \max\left\{-\log\left((1 - \epsilon_N) \frac{p_N}{\mu}\right), 0\right\} + p_N \left(-\log\left(\frac{1}{N}\right)\right) \frac{\sigma^2}{\frac{\epsilon_N^2}{(1 - \epsilon_N)} Np_N \mu} && \text{(EC.20)}
\end{aligned}$$

$$= p_N \zeta_N + p_N \max\left\{\log \frac{\mu}{(1 - \epsilon_N) p_N}, 0\right\} + \frac{\sigma^2 \log N}{\frac{\epsilon_N^2}{(1 - \epsilon_N)} \mu N}, \quad \text{(EC.21)}$$

where the first equality follows from (EC.13), the second equality follows from the facts that $\tau_N \leq N$ and $r_N = N - \tau_N + 1$ when $\tau_N \geq 1$, and the third equality follows from the fact that $r_N \geq 1$. Eq. (EC.20) follows by the monotonicity of $\log(\cdot)$ and by applying Eq. (EC.19).

Now setting $\epsilon_N = N^{-\frac{1}{3}}$ and substituting in (EC.21) gives

$$\mathcal{R}_{TB}(p_N) \leq p_N \zeta_N + p_N \max\left\{\log \frac{\mu}{(1 - N^{-\frac{1}{3}}) p_N}, 0\right\} + \frac{(1 - N^{-\frac{1}{3}}) \sigma^2 \log N}{\mu N^{\frac{1}{3}}}. \quad \text{(EC.22)}$$

Maximizing Eq. (EC.22) over p_N gives $\bar{p}_N := \mu / \exp(1 + \log(1 - N^{-\frac{1}{3}}) - \zeta_N)$. Plugging in \bar{p}_N into Eq. (EC.22) gives the desired upper bound, and combining with Eq. (EC.18) yields

$$\frac{\mu}{e} \log \left(\frac{1}{\frac{1}{e} + \frac{1 + \frac{\sigma^2}{\mu^2}}{N}} \right) \leq \mathcal{R}_{TB} \leq \frac{\mu}{e^{1-\zeta_N} (1 - N^{-\frac{1}{3}})} + \frac{(1 - N^{-\frac{1}{3}}) \sigma^2 \log N}{\mu N^{\frac{1}{3}}}. \quad (\text{EC.23})$$

Taking the limit of both sides of Eq. (EC.23) completes the proof. \square

Proof of Theorem 4.

(a) *Unique Box.* Suppose the optimal price p_N^* does not converge to μ , i.e., there exist $\epsilon > 0$ such that $|p_N^* - \mu| > \epsilon$ infinitely often. First, consider the case where $p_N^* < \mu - \epsilon$ infinitely often. Since the normalized selling volume is at most 1, then the revenue must be less than $\mu - \epsilon$ infinitely often. However, this contradicts the fact that \mathcal{R}_{UB} converges to μ .

Now consider the case that $p_N^* > \mu + \epsilon$ infinitely often. Using Chebyshev's inequality, the probability that a customer purchases the first loot box is at most

$$\mathbb{P} \left(\frac{\sum_i V_i}{N} \geq p \right) \leq \frac{\sigma^2}{(p_N^* - \mu)^2 N}.$$

Thus, $\frac{\sigma^2}{(p_N^* - \mu)^2 N}$ is also an upper bound on the normalized sales volume, since the best case the customer buys the maximum N unique boxes. Thus, an upper bound on the the normalized revenue when $p_N^* > \mu + \epsilon$ is $\frac{\sigma^2 p_N^*}{(p_N^* - \mu)^2 N}$. Note that $\frac{\sigma^2 p_N^*}{(p_N^* - \mu)^2 N}$ is decreasing p_N^* when $p_N^* > \mu + \epsilon$, so an even greater upper bound on the revenue in this case is $\frac{\sigma^2 (\mu + \epsilon)}{\epsilon^2 N}$. Since this upper bound tends to 0 as N tends to ∞ , then this contradicts the fact that \mathcal{R}_{UB} converges to μ and thus p_N^* cannot be greater than $\mu + \epsilon$ infinitely often.

Now we consider the expected fraction of unique items collected by the customer, which is also the expected normalized selling volume for the unique box strategy. Since the normalized selling volume is upper bounded by 1, if it does not converge to 1, \mathcal{R}_{UB} cannot converge to μ given that the optimal price converges to μ . Hence the expected selling volume converges to 1.

Finally, since the expected customer valuation is μ and \mathcal{R}_{UB} converges to μ , then no utility is left for the customer and therefore the normalized customer surplus converges to 0.

(b) *Traditional Box.* We first show that the optimal price converges to $\frac{\mu}{e}$. Suppose the optimal price p_N^* does not converge to $\frac{\mu}{e}$, i.e., there exists $\epsilon > 0$ such that $|p_N^* - \frac{\mu}{e}| > \epsilon$ infinitely often. Recall from Eq. (EC.22) that the revenue by using any price p_N is upper bounded by

$$p_N \zeta_N + p_N \max \left\{ \log \frac{\mu}{(1 - N^{-\frac{1}{3}}) p_N}, 0 \right\} + \frac{(1 - N^{-\frac{1}{3}}) \sigma^2 \log N}{\mu N^{\frac{1}{3}}},$$

and this upper bound converges to $p_N \max \{ \log \frac{\mu}{p_N}, 0 \}$. Note that $p_N \max \{ \log \frac{\mu}{p_N}, 0 \} < \frac{\mu}{e}$ for any $p_N \neq \frac{\mu}{e}$. Therefore, using a price bounded away from $\frac{\mu}{e}$ infinitely often results in a revenue that is bounded away $\frac{\mu}{e}$ infinitely often. This contradicts the fact that \mathcal{R}_{TB} converges to $\frac{\mu}{e}$ and thus p_N^* cannot be bounded away from $\frac{\mu}{e}$ infinitely often.

The fraction of unique items collected by the customer is given by $\tau_N(p)/N = (N + 1 - r_N(p)) \mathbb{I}_{r_N(p) \leq N} / N$. By Eq. (EC.14), $\mathbb{E}[r_N(p)]/N$ converges to p/μ . When $p < \mu$, by Theorem 7.1 in (Gut 2009), $\mathbb{E}[r_N(p)]/N$ is uniformly integrable and $\mathbb{E}[r_N(p) \mathbb{I}_{r_N(p) \leq N}]/N$ converges to p/μ . Plugging in $\lim_{N \rightarrow \infty} p_N^* = \mu/e$, we have

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E}[\tau_N(p_N^*)]}{N} = \lim_{N \rightarrow \infty} \frac{\mathbb{E}[(N - r_N(p_N^*) + 1) \mathbb{I}_{r_N(p_N^*) \leq N}]}{N} = 1 - \lim_{N \rightarrow \infty} \frac{\mathbb{E}[r_N(p_N^*) \mathbb{I}_{r_N(p_N^*) \leq N}]}{N} = 1 - \frac{\mu/e}{\mu} = 1 - \frac{1}{e}.$$

For the selling volume, note that for any p , Eq. (EC.17) and Eq. (EC.22) implies that

$$\lim_{N \rightarrow \infty} \frac{\mathcal{R}_{TB}(p)}{p} = \max(0, \log \frac{\mu}{p}). \quad (\text{EC.24})$$

Plugging in $p^* = \mu/e$ gives the normalized selling volume in the limit, which is 1.

Finally, the customer surplus is the total utility from the unique items $\sum_{i=1}^{\tau N} V_i$ minus the total price paid. Hence we have

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathbb{E}[\text{Normalized Surplus}] &= \lim_{N \rightarrow \infty} \frac{\mathbb{E}[\sum_{i=1}^{\tau N} V_i]}{N} - \lim_{N \rightarrow \infty} \mathcal{R}_{TB} \\ &= \lim_{N \rightarrow \infty} \frac{\mathbb{E}[\tau_N(p_N^*)] \mathbb{E}[V_i]}{N} - \lim_{N \rightarrow \infty} \mathcal{R}_{TB} \\ &= (1 - \frac{1}{e})\mu - \frac{\mu}{e} \\ &= (1 - \frac{2}{e})\mu. \end{aligned}$$

□

A.3. Omitted Proofs from Section 4

Proof of Proposition 2. For unique boxes, we show that the revenue from the size-1 case is dominated by the revenue of the size- j case with a simple coupling argument. Since the asymptotic revenue in the size-1 case is μ by Theorem 2, this implies that the asymptotic revenue for the size- j case is also μ . Now suppose p is the price in the size-1 case and set jp to be the price in the size- j case. If a customer bought τ loot boxes in the size-1 case and would like to buy the next size-1 box given that they owned the set S_τ , then we claim that the same customer would have bought a size- j box. This follows from the fact that the valuation of a size- j unique box in this state is exactly j times the valuation of a size-1 box. Thus, since the price is also scaled by j , the decision of purchasing a loot box in period τ is perfectly coupled, which concludes the proof.

For traditional boxes, we use a more complex coupling argument to show that the revenue from the size-1 case is very close to the revenue of the size- j case with a price slightly lower than jp . Let p be the price of the size-1 box and let $pN(1 - (1 - 1/N)^j)$ be the price of the size- j box. If a customer has purchased τ size-1 box with inventory state S_τ , and would like to buy the next size-1 box, then we claim that the same customer would like to buy a size- j box given the same situation. This follows from the fact that the customer may get a specific item with probability $1 - (1 - 1/N)^j$, and the valuation of a size- j unique box after owning S_τ is $(1 - (1 - 1/N)^j) \sum_{[N] \setminus S_\tau} V_i$, while the corresponding valuation of the size-1 box is $\frac{1}{N} \sum_{[N] \setminus S_\tau} V_i$. Therefore, a size- j box is purchased in period τ if and only if a size-1 box would have been purchased:

$$\frac{1}{N} \sum_{[N] \setminus S_\tau} V_i \geq p \iff (1 - (1 - 1/N)^j) \sum_{[N] \setminus S_\tau} V_i \geq pN(1 - (1 - 1/N)^j).$$

Hence, if a customer stops after purchasing τ size-1 boxes, along with the same sampling path he will stop after purchasing $\lceil \tau/j \rceil$ size- j boxes. Note that $jp \geq pN(1 - (1 - 1/N)^j)$, so the normalized revenue generated by size- j box is bounded as

$$\frac{pN(1 - (1 - 1/N)^j)}{jp} \mathcal{R}_{TB}(p) \leq \mathcal{R}_{TB}^j(pN(1 - (1 - 1/N)^j)) \leq \frac{pN(1 - (1 - 1/N)^j)}{jp} \mathcal{R}_{TB}(p) + \frac{pN(1 - (1 - 1/N)^j)}{N}.$$

Taking the limit of the above as $N \rightarrow \infty$ leads to $\lim_{N \rightarrow \infty} \mathcal{R}_{TB}^j(jp) = \lim_{N \rightarrow \infty} \mathcal{R}_{TB}(p)$. Since the optimal TB revenue is $\frac{\mu}{e}$ by Theorem 3, this concludes the proof. □

Proof of Proposition 3. In this proposition we show that unique box with uniform allocation is asymptotically optimal. We modify the random walk X_t in the proof of Theorem 2 into a stochastic process $\{X'_t, t \geq 0\}$. For $t \leq N$, let X'_t be the net utility of a random customer after opening t unique boxes. For $t > N$, simply let $X'_t - X'_{t-1} = \bar{\mu} - p$. Note X'_N has mean $N(\bar{\mu} - p)$ and variance $N\bar{\sigma}^2$, where $\bar{\sigma}^2 = \sum_{m=1}^M \beta_m \sigma_m^2$. Also, the expectation of $X'_t - X'_{t-1}$ is $\bar{\mu} - p$ for any $t \geq 1$. Note $\{X'_t\}_{t=1}^\infty$ is not a stationary random walk, since its step lengths are correlated. However, $\{X'_t\}_{t=1}^\infty$ satisfies the Markovian property, as for every t , X'_{t+1} depends only on the number of items in each class that are not yet owned. Hence, following the proof of Theorem 2, the Wald's equations (Eq. (EC.6)) and Chebyshev's inequalities (Eq. (EC.8), Eq. (EC.11)) are still valid. The only difference is the overshoot term, $\mathbb{E}[X'_{\bar{\tau}_N} - (1 - k_N)N\epsilon_N]$. By Theorem 2 in (Lorden 1970), it is bounded by $(\sigma_{\max}^2 + \mu_{\max}^2)/\epsilon_N + ((\sigma_{\max}^2 + \mu_{\max}^2)(1 - k_N)\epsilon_N N/\epsilon_N)^{1/2}$, which will not influence the limit and the convergence rate. Thus, the limiting result remains the same. \square

Proof of Proposition 4. Fix $k \in (0, 1)$, a probability vector \mathbf{d} , and let $p = \sum_{m=1}^M d_m \mu_m e^{-\frac{d_m}{\beta_m} k}$. We shall show that the normalized number of loot box purchases made by a customer under the pricing and allocation strategy (p, \mathbf{d}) , $\mathbb{E}[Q^N(p, \mathbf{d})]$, tends to k as $N \rightarrow \infty$. For clarity, we prove the lower and upper bounds separately.

Lower Bound: $\lim_{N \rightarrow \infty} \mathbb{E}[Q^N(p, \mathbf{d})] \geq k$

Given $p = \sum_{m=1}^M d_m \mu_m e^{-\frac{d_m}{\beta_m} k}$, we first bound the probability that $Q^N(p, \mathbf{d}) < (1 - \epsilon)k$. Note since a customer's valuation for the next loot box decreases monotonically after each purchase, the event (a) $Q^N(p, \mathbf{d}) \leq (1 - \epsilon)k$ is equivalent to the event (b) the customer's valuation for the loot box is less than p after they have opened $(1 - \epsilon)kN$ boxes. We will bound this event by applying Chebyshev's inequality, for which we will need estimates of both the mean and variance of customer's valuation after opening $(1 - \epsilon)kN$ boxes. Let Z_i^m be an indicator random variable taking value 1 if item i from class m has not been revealed after $(1 - \epsilon)kN$ purchases, and 0 otherwise. When the class is clear from the context we will drop the superscript. Now, after each purchase the probability that item i in class m is obtained is $\frac{d_m}{\beta_m N}$, thus the expectation of Z_i^m is,

$$\mathbb{E}[Z_i^m] = \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN}.$$

Recall that G_m denotes the set of items in class m . For a random customer, since V_i and Z_i^m are independent, the valuation of the next loot box after $(1 - \epsilon)kN$ purchases is given by $\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i^m$, and the expected valuation for a loot box after $(1 - \epsilon)kN$ purchases is,

$$\mathbb{E}\left[\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i^m\right] = \sum_{m=1}^M \sum_{i \in G_m} \mu_m \frac{d_m}{\beta_m N} \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN} = \sum_{m=1}^M \mu_m d_m \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN}.$$

Moreover, observe that the set of indicators $\{Z_i^m\}_{i=1}^{|G_m|}$ is negatively correlated for all i and m since for any two different items, if one is not revealed so far then the other is more likely to have been revealed. Thus the variance can be bounded by,

$$\text{Var}\left(\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i^m\right)$$

$$\begin{aligned}
&\leq \sum_{m=1}^M \sum_{i \in G_m} \text{Var} \left(\frac{d_m}{\beta_m N} V_i Z_i \right) && (\{Z_i^m\} \text{ Neg. Corr.}) \\
&= \sum_{m=1}^M \sum_{i \in G_m} \left(\frac{d_m}{\beta_m N} \right)^2 \left(\mathbb{E} [V_i^2 Z_i^2] - (\mathbb{E}[V_i Z_i])^2 \right) \\
&= \sum_{m=1}^M \sum_{i \in G_m} \left(\frac{d_m}{\beta_m N} \right)^2 \left(\mathbb{E} [V_i^2 Z_i] - (\mathbb{E}[V_i Z_i])^2 \right) && (Z_i^m \in \{0, 1\}) \\
&= \sum_{m=1}^M \sum_{i \in G_m} \left(\frac{d_m}{\beta_m N} \right)^2 \left((\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN} - \mu_m^2 \left(1 - \frac{d_m}{\beta_m N}\right)^{2(1-\epsilon)kN} \right) && (V_i, Z_i \text{ independent}) \\
&\leq \sum_{m=1}^M \sum_{i \in G_m} \left(\frac{d_m}{\beta_m N} \right)^2 (\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN} \\
&= \sum_{m=1}^M \frac{d_m^2}{\beta_m N} (\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN} && \text{(EC.25)}
\end{aligned}$$

Now applying Chebyshev's Inequality to the event that the mean valuation is less than p after opening $(1 - \epsilon)kN$ boxes, we have

$$\begin{aligned}
&\mathbb{P} (Q^N(p, \mathbf{d}) \leq (1 - \epsilon)k) \\
&= \mathbb{P} \left(\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i < p \right) && \text{(Event (a), (b) equivalent)} \\
&= \mathbb{P} \left(\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i - \mathbb{E} \left[\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i \right] < p - \mathbb{E} \left[\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i \right] \right) \\
&\leq \mathbb{P} \left(\left| \sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i - \mathbb{E} \left[\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i \right] \right| > \mathbb{E} \left[\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i \right] - p \right) \\
&\leq \frac{\text{Var} \left(\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i \right)}{\left(\mathbb{E} \left[\sum_{m=1}^M \sum_{i \in G_m} \frac{d_m}{\beta_m N} V_i Z_i \right] - p \right)^2} && \text{(Chebyshev's Inequality)} \\
&\leq \frac{\sum_{m=1}^M \frac{d_m^2}{\beta_m N} (\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN}}{\left(\sum_{m=1}^M \mu_m d_m \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN} - p \right)^2} && \text{(From Eq. (EC.25))} \\
&= \sum_{m=1}^M \frac{d_m^2 (\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN}}{\beta_m N \left(\sum_{m=1}^M \mu_m d_m \left(\left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN} - e^{-\frac{d_m}{\beta_m} k} \right) \right)^2} && \text{(EC.26)}
\end{aligned}$$

Taking the limit as N tends to infinity, the numerator of Eq. (EC.26) approaches a constant $d_m^2 (\mu_m^2 + \sigma_m^2) e^{-\frac{d_m}{\beta_m} k (1-\epsilon)}$. For the denominator, the term $\left(1 - \frac{d_m}{\beta_m N}\right)^{(1-\epsilon)kN} - e^{-\frac{d_m}{\beta_m} k}$ goes to $e^{-\frac{d_m}{\beta_m} k (1-\epsilon)} - e^{-\frac{d_m}{\beta_m} k}$, which is a constant, so the denominator goes to infinity. Thus for any $\epsilon > 0$, there exists C_1 such that for all $N > C_1$, $\mathbb{P} (Q^N(p, \mathbf{d}) \leq (1 - \epsilon/2k)k) \leq \epsilon/2k$. Applying Eq. (EC.26) yields a lower bound of $\mathbb{E} [Q^N(p, \mathbf{d})]$,

$$\begin{aligned}
\mathbb{E} [Q^N(p, \mathbf{d})] &= \mathbb{E} [Q^N(p, \mathbf{d}) \mathbb{I}_{Q^N(p, \mathbf{d}) \leq (1-\epsilon/2k)k}] + \mathbb{E} [Q^N(p, \mathbf{d}) \mathbb{I}_{Q^N(p, \mathbf{d}) > (1-\epsilon/2k)k}] \\
&\geq 0 + \left(1 - \frac{\epsilon}{2k}\right) k \left(1 - \frac{\epsilon}{2k}\right) \\
&= k \left(1 - \frac{\epsilon}{2k}\right)^2
\end{aligned}$$

$$\geq k\left(1 - \frac{\epsilon}{k}\right) = k - \epsilon, \quad (\text{EC.27})$$

which implies that the lower bound converges to k as ϵ goes to 0.

Upper Bound: $\lim_{N \rightarrow \infty} \mathbb{E}[Q^N(p, \mathbf{d})] \leq k$

Similar to the lower bound, we first control the event that $Q^N(p, \mathbf{d}) > (1 + \epsilon)k$. Let $Z_i^m = 1$ now denote the event that after opening $(1 + \epsilon)kN$ loot boxes, item i in group m is still not revealed. As before we will omit the superscript when it is clear from context. Following the derivation of (EC.26), we may bound the probability of this event by

$$\mathbb{P}(Q^N(p, \mathbf{d}) > (1 + \epsilon)k) \leq \sum_{m=1}^M \frac{d_m^2 (\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1+\epsilon)kN}}{\beta_m N \left(\sum_{m=1}^M \mu_m d_m \left(e^{-\frac{d_m}{\beta_m} k} - \left(1 - \frac{d_m}{\beta_m N}\right)^{(1+\epsilon)kN}\right)\right)^2}. \quad (\text{EC.28})$$

Now we will choose $\epsilon = -\log(1 - N^{-1/3})/k$. Substituting our choice of ϵ into the denominator of Eq. (EC.28) we may obtain a lower bound,

$$\begin{aligned} & \beta_m N \left(\sum_{m=1}^M \mu_m d_m \left(e^{-\frac{d_m}{\beta_m} k} - \left(1 - \frac{d_m}{\beta_m N}\right)^{(1+\epsilon)kN}\right)\right)^2 \\ & \geq \beta_m \left(\sum_{m=1}^M \mu_m d_m N^{1/2} \left(e^{-\frac{d_m}{\beta_m} k} - e^{-\frac{d_m}{\beta_m} (1+\epsilon)k}\right)\right)^2 && \text{(Taylor expansion of } e^x \text{)} \\ & = \beta_m \left(\sum_{m=1}^M \mu_m d_m e^{-\frac{d_m}{\beta_m} k} N^{1/2} \left(1 - \left(1 - N^{-1/3}\right)^{\frac{d_m}{\beta_m}}\right)\right)^2 \\ & \geq \beta_m \left(\sum_{m=1}^M \mu_m d_m e^{-\frac{d_m}{\beta_m} k} N^{1/2} \left(1 - \frac{1}{1 + \frac{d_m}{\beta_m} N^{-1/3}}\right)\right)^2 && \text{(Bernoulli's Inequality)} \\ & = \beta_m \left(\sum_{m=1}^M \mu_m d_m e^{-\frac{d_m}{\beta_m} k} \frac{\frac{d_m}{\beta_m} N^{1/2}}{N^{1/3} + \frac{d_m}{\beta_m}}\right)^2. \end{aligned}$$

Plugging back into Eq. (EC.28), the probability that customer purchases more than $(1 - \log(1 - N^{-1/3}))N$ boxes is then bounded above by,

$$\mathbb{P}(Q^N(p, \mathbf{d}) > (1 - \log(1 - N^{-1/3}))N) \leq \sum_{m=1}^M \frac{d_m^2 (\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1 - \log(1 - N^{-1/3}))/k kN}}{\beta_m \left(\sum_{m=1}^M \mu_m d_m e^{-\frac{d_m}{\beta_m} k} \frac{\frac{d_m}{\beta_m} N^{1/2}}{N^{1/3} + \frac{d_m}{\beta_m}}\right)^2}.$$

Finally, returning to $Q_{\mathbf{d}(p)}^N$, a trivial upper bound on $\mathbb{E}[Q^N(p, \mathbf{d})]$ is given by

$$\begin{aligned} \mathbb{E}[Q^N(p, \mathbf{d})] & \leq \mathbb{E}[\# \text{ of purchases to collect all the items}] \\ & \leq \sum_{m=1}^M \mathbb{E}[\# \text{ of purchases to collect all the items in class } m] \\ & = \sum_{m=1}^M \mathbb{E}\left[\text{Geo}(d_m) + \text{Geo}\left(\frac{d_m(\beta_m N - 1)}{\beta_m N}\right) + \dots + \text{Geo}\left(\frac{d_m}{\beta_m N}\right)\right] \\ & = \sum_{m=1}^M \frac{\beta_m N}{d_m \beta_m N} + \frac{\beta_m N}{d_m(\beta_m N - 1)} + \dots + \frac{\beta_m N}{d_m} \end{aligned}$$

$$\begin{aligned}
&= \sum_{m=1}^M \frac{\beta_m N}{d_m} \left(\frac{1}{\beta_m N} + \frac{1}{\beta_m N - 1} + \cdots + 1 \right) \\
&\leq \sum_{m=1}^M \frac{\beta_m N}{d_m} (\log(\beta_m N) + 1) \\
&\leq \sum_{m=1}^M \frac{\beta_m N}{d_m} (\log N + 1).
\end{aligned}$$

Thus the expected number of purchases is upper bounded by $\mathbb{E}[Q^N(p, \mathbf{d})] \leq \sum_{m=1}^M \frac{\beta_m N}{d_m} (\log N + 1)$ for any price p . Now we can build an upper bound:

$$\begin{aligned}
\mathbb{E}[Q^N(p, \mathbf{d})] &= \mathbb{E}[Q^N(p, \mathbf{d}) \mathbb{I}_{Q^N(p, \mathbf{d}) \leq (1+\epsilon)kN}] + \mathbb{E}[Q^N(p, \mathbf{d}) \mathbb{I}_{Q^N(p, \mathbf{d}) > (1+\epsilon)kN}] \\
&\leq (1 - \log(1 - N^{-1/3})/k) kN + N \sum_{m=1}^M \frac{\beta_m}{d_m} (\log N + 1) \sum_{m=1}^M \frac{d_m^2 (\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1 - \log(1 - N^{-1/3})/k)kN}}{\beta_m \left(\sum_{m=1}^M \mu_m d_m e^{-\frac{d_m}{\beta_m} k} \frac{\frac{d_m}{\beta_m} N^{1/2}}{N^{1/3} + \frac{d_m}{\beta_m}}\right)^2}
\end{aligned}$$

Taking $N \rightarrow \infty$ on both sides, we have

$$\begin{aligned}
&\lim_{N \rightarrow \infty} \frac{\mathbb{E}[Q^N(p, \mathbf{d})]}{N} \\
&\leq \lim_{N \rightarrow \infty} (1 - \log(1 - N^{-1/3})/k) k + \lim_{N \rightarrow \infty} \sum_{m=1}^M \frac{\beta_m}{d_m} (\log N + 1) \sum_{m=1}^M \frac{d_m^2 (\mu_m^2 + \sigma_m^2) \left(1 - \frac{d_m}{\beta_m N}\right)^{(1 - \log(1 - N^{-1/3})/k)kN}}{\beta_m \left(\sum_{m=1}^M \mu_m d_m e^{-\frac{d_m}{\beta_m} k} \frac{\frac{d_m}{\beta_m} N^{1/2}}{N^{1/3} + \frac{d_m}{\beta_m}}\right)^2} \\
&= k + \lim_{N \rightarrow \infty} \sum_{m=1}^M \frac{d_m^2 (\mu_m^2 + \sigma_m^2) e^{-\frac{d_m}{\beta_m} k} \sum_{m=1}^M \frac{\beta_m}{d_m} (\log N + 1)}{\beta_m \left(\sum_{m=1}^M \mu_m d_m e^{-\frac{d_m}{\beta_m} k} \frac{\frac{d_m}{\beta_m} N^{1/2}}{N^{1/3} + \frac{d_m}{\beta_m}}\right)^2} \\
&= k + \lim_{N \rightarrow \infty} \sum_{m=1}^M \frac{d_m^2 (\mu_m^2 + \sigma_m^2) e^{-\frac{d_m}{\beta_m} k} \sum_{m=1}^M \frac{\beta_m}{d_m} \left(1 + \frac{1}{\log N}\right)}{\beta_m \left(\sum_{m=1}^M \mu_m d_m e^{-\frac{d_m}{\beta_m} k} \frac{\frac{d_m}{\beta_m} N^{1/2}}{(\log N)^{1/2} (N^{1/3} + \frac{d_m}{\beta_m})}\right)^2} \\
&= k. \tag{EC.29}
\end{aligned}$$

Thus combining Eq. (EC.27) and Eq. (EC.29), we have

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E}[Q^N(p, \mathbf{d})]}{N} = k.$$

□

Proof of Proposition 5.

(a) Let $V_i^c := \max\{V_i, c\}$ be the modified valuation of an item that has salvage value c , and let F_c denote the distribution of V_i^c . Let $\tilde{\eta}$ be the mean of V_i^c . By Theorem 2 and Theorem 4, as the number of items $N \rightarrow \infty$, the optimal price tends to $\tilde{\eta}$ and the expected proportion of items obtained tends to 1. Since all items are obtained in expectation, the proportion of items salvaged tends to $F(c)$. Thus the normalized cost of salvages by the customer is $\lim_{N \rightarrow \infty} \frac{E[\#\text{ Items Salvaged}]}{N} c = F(c)c$. Together, the normalized revenue is then $\tilde{\eta} - F(c)c$. Noting $\tilde{\eta}$ can be rewritten as $\tilde{\eta} = E[\max\{V_i, c\}] = F(c)c + \overline{F}(c)\mathbb{E}[V_i | V_i > c] = F(c)c + \overline{F}(c)\eta$, then the normalized revenue becomes $E[\max\{V_i, c\}] - F(c)c = \overline{F}(c)\eta$.

For customer surplus, note that the customer all items in the limit, garnering expected utility of $\tilde{\eta}$. The cost to the customer is the revenue $\bar{F}(c)\eta$, so the customer surplus is $\tilde{\eta} - \bar{F}(c)\eta = F(c)c$.

(b) Following Theorem 3, we consider a modified random walk for customers of a traditional box strategy with salvage cost c . Let $Y'_j = \sum_{i=j+1}^N V_i^c + jc$. For a random customer, Y'_j/N is the expected valuation of the traditional box after receiving j unique items. If the new item is at value greater than c , then Y'_j is decreased by $V_{j+1} - c$, otherwise it decreases by 0. Hence, the mean step length is given by

$$\mathbb{E}[Y'_{j+1} - Y'_j] = \bar{F}(c)(\eta - c) + F(c) \cdot 0 = \bar{F}(c)(\eta - c).$$

Also, note that $Y'_N = Nc$. The random walk Y'_j is still a decreasing process, which means that the valuation of the box is decreasing as customers collect more and more new items. A customer purchases until the first time such that $\frac{Y'_t}{N} < p \iff Y'_t < Np \iff Y'_t - Nc < N(p - c)$. Now consider the random walk $\{Y'_j - Nc\}$, which ends with 0, and is weakly decreasing with mean jump length $\bar{F}(c)(\eta - c)$. Let $\tau(p)$ be the first passage time of $\{Y'_j - Nc\}$ hitting the line $N(p - c)$ from above. The problem of approximating $\tau(p)$ is exactly the same problem of approximating the expected selling volume of a vanilla traditional box in Theorem 3 with mean $\bar{F}(c)(\eta - c)$ and price $p - c$. Recall that in the proof of Theorem 4, we show that the limiting selling volume for a vanilla traditional box is $\max(\log \frac{\mu}{p}, 0)$ (see (EC.24)). So in the case with salvage c , we know that the normalized selling volume is given by

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E}[\text{selling volume}]}{N} = \max\left(\log \frac{\bar{F}(c)(\eta - c)}{p - c}, 0\right),$$

and for the nontrivial case $p - c \leq \bar{F}(c)(\eta - c)$, the selling volume is simply $\log \frac{\bar{F}(c)(\eta - c)}{p - c}$. The net revenue is the revenue subtracted by the salvage cost. Note that only the new items with value greater than c are not salvaged. The number of unique items is $\tau(p)$, and by the discussion in Theorem 4, $\tau(p)/N$ converges to $1 - p/\mu$, which is $1 - \frac{p - c}{\bar{F}(c)(\eta - c)}$ in the new problem. Hence the limiting revenue with price $p \leq c + \bar{F}(c)(\eta - c)$ is

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathcal{R}_{TB}(c, p) &= \lim_{N \rightarrow \infty} p \frac{\mathbb{E}[\text{selling volume}]}{N} - c \frac{\mathbb{E}[\text{selling volume}]}{N} - \lim_{N \rightarrow \infty} \frac{\mathbb{E}[\# \text{ of unique item with value } > c]}{N} \\ &= \lim_{N \rightarrow \infty} (p - c) \frac{\mathbb{E}[\text{selling volume}]}{N} + \lim_{N \rightarrow \infty} c \frac{\bar{F}(c)\mathbb{E}[\tau(p)]}{N} \\ &= (p - c) \log \frac{\bar{F}(c)(\eta - c)}{p - c} + c(\bar{F}(c) - \frac{p - c}{\eta - c}). \end{aligned}$$

Maximizing over the price yields $p = c + e^{-\frac{\eta}{\eta - c}} \bar{F}(c)(\eta - c)$. Plugging in p gives our desired revenue $\bar{F}(c)(\eta - c) \left(\frac{c}{\eta - c} + e^{-\frac{\eta}{\eta - c}}\right)$.

Finally, customer surplus is the total utility from the unique items that the customer keeps minus the total cost (i.e. revenue of the seller). Hence we have

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathbb{E}[\text{Normalized Surplus}] &= \lim_{N \rightarrow \infty} \frac{\mathbb{E}[\text{Utility from unique items with value } > c]}{N} - \lim_{N \rightarrow \infty} \mathcal{R}_{TB}^c \\ &= \lim_{N \rightarrow \infty} \frac{\eta \mathbb{E}[\# \text{ of unique items with value } > c]}{N} - \lim_{N \rightarrow \infty} \mathcal{R}_{TB}^c \\ &= \eta \left(\bar{F}(c) - \frac{p - c}{\eta - c} \right) - \bar{F}(c)(\eta - c) \left(\frac{c}{\eta - c} + e^{-\frac{\eta}{\eta - c}} \right) \\ &= \bar{F}(c) \left((\eta - c) - (2\eta - c)e^{-\frac{\eta}{\eta - c}} \right). \quad \square \end{aligned}$$

Appendix B: Omitted Examples

EXAMPLE EC.1 (UNIFORM UNIQUE BOX MAY NOT BE OPTIMAL). Consider a unique box with two items facing a customer with non-uniform allocation probabilities. At the time of the first purchase, the customer has probability q to receive item 1, and $1 - q$ to receive the item 2. Since it is a unique box, upon second purchase the customer will receive the unowned item with probability 1. Now suppose items are valued as either 0 or 1, with probability 0.5.

Consider a unique box with certain allocations, $q = 1$. In this case, the first purchase of a loot box always yields item 1, and the the second purchase then gives the remaining item 2. For this box, the optimal price is 1, and the selling volume would be $\frac{1}{4}(1 + 2)$, i.e., customers whose valuation is (1,0) will buy 1 box, and those with valuation (1,1) will buy 2 boxes. The resulting revenue is $\frac{3}{4}$. On the other hand, if $q = 0.5$ then we have a uniform unique box, and the corresponding optimal price can be checked to be 0.5. The the uniform allocation and corresponding price induce selling volume $\frac{1}{4}(1.5 + 1.5 + 2)$, i.e., customers whose valuation is (0,1) or (1,0) will buy 1.5 boxes on average, and customers whose valuation is (1,1) will buy 2 boxes. The resulting revenue is only $\frac{5}{8}$. \square

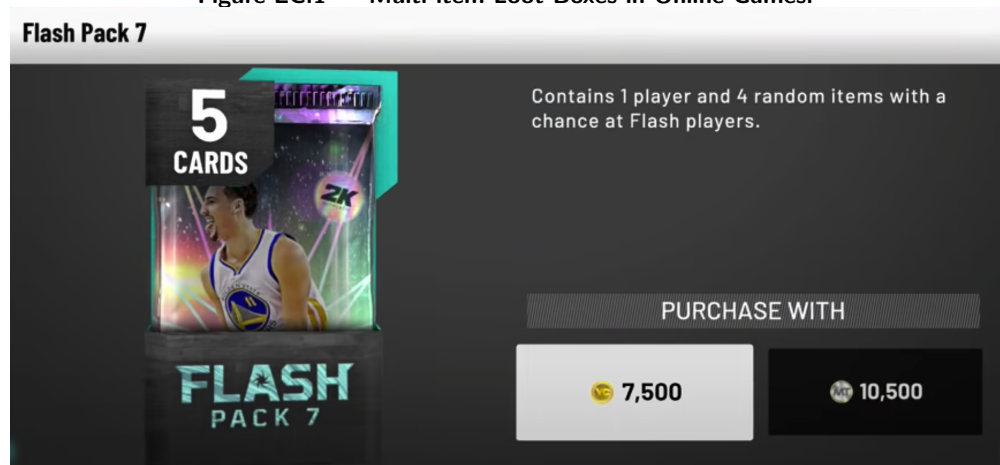
EXAMPLE EC.2 (NO K-STEP LOOK-AHEAD POLICY IS OPTIMAL FOR CUSTOMERS). Let the number of products be N and the price of each unique box be 2.5. Now consider a customer whose realized valuations for the products are $(N, \frac{N}{N-1}, \dots, \frac{N}{N-1})$. If the customer is myopic, then they will not buy a single unique box since the expected utility of the first loot box is $\frac{N+(N-1)\frac{N}{N-1}}{N} - 2.5 = -0.5 < 0$.

Now consider the following policy: purchase unique boxes until you obtain the item which is valued at N . In expectation such a strategy requires $\frac{N}{2}$ purchases and yields utility $N + \frac{N}{2} - 1 - \frac{2.5N}{2} = \frac{N}{4} - 1$. It is straight forward to show this policy is optimal given these parameters. Further, the expected utility for each of the first $\approx \frac{N}{3}$ loot box purchases before obtaining the high valued item is negative since, if a customer has acquired q percent of the catalog without obtaining the high valued item, their expected utility for the next box is $\frac{N+((1-q)N-1)\frac{N}{N-1}}{(1-q)N} - 2.5$ which is less than zero when $q \leq \frac{1}{3}$. Thus no policy that considers only a fixed number of future purchases can be optimal. \square

Appendix C: Omitted Figures

In this section we include additional figures depicting various forms of loot box selling in practice.

Figure EC.1 Multi-item Loot Boxes in Online Games.



Note. Depicted is a multi-item loot box in the game *NBA 2K20*. Each box contains 5 cards.

Figure EC.2 Loot Box with Multiple Classes.

Crate Odds		
?	Mythic Item	0.3%
?	Legendary Item	1.0%
?	Epic Item	22.7%
?	Rare Item	76%

Note. In the game *PlayerUnknown's Battlegrounds*, the traditional box contains four classes of items: Mythic, Legendary, Epic, and Rare. The allocation probability varies across classes, however items within the same class have the same probability.

Figure EC.3 Salvage System in Dota 2.



Note. In the game *Dota 2*, players can trade in 6 unwanted items for a new loot box plus 2000 shards, a form of in-game currency.