Proceedings of the Sixteenth International AAAI Conference on Web and Social Media (ICWSM 2022)

Quantifying the Creator Economy: A Large-Scale Analysis of Patreon

Lana El Sanyoura and Ashton Anderson
University of Toronto
lana.elsanyoura@mail.utoronto.ca, ashton@cs.toronto.edu

Abstract

In recent years, the "creator economy" has emerged as a disruptive force in creative industries. Independent creators can now reach large and diverse audiences through online platforms, and membership platforms have emerged to connect these creators with fans who are willing to financially support them. However, the structure and dynamics of how membership platforms function on a large scale remain poorly understood. In this work, we develop an analysis framework for the study of membership platforms and apply it to the complete set of Patreon pledges exceeding $2 billion since its inception in 2013 until the end of 2020. We analyze Patreon activity through three perspectives: patrons (demand), creators (supply), and the platform as a whole. We find several important phenomena that help explain how membership platforms operate. Patrons who pledge to a narrow set of creators are more loyal, but churn off the platform more often. High-earning creators attract large audiences, but these audiences are less likely to pledge to other creators. Over its history, Patreon diversified into many topics and launched higher-earning creators over time. Our analysis framework and results shed light on the functioning of membership platforms and have implications for the creator economy.

Introduction

One of the most significant recent developments in the online world has been the rise of the "creator economy". Independent creators spanning industries from video entertainment and podcasting to art and music can now reach large and diverse audiences through online platforms, without the help of agencies or industry representatives. *Membership platforms* have emerged to connect these creators with fans who are willing to financially support them. Patreon, the largest membership platform, connects millions of creators with millions of patrons and processes over $1 billion in payments to creators annually. Patreon currently processes more than $100M to creators every month and has directed more than $2B in revenue to creators since its inception in 2013. As such, it is by far the largest membership platform in the world, and has played a leading role in the development of the creator economy. Since Patreon's introduction of creator-fan subscriptions in 2013, other membership platforms have emerged, such as BuyMeaCoffee, OnlyFans, and Ko-fi. This membership model has become so popular that existing social media platforms have started to incorporate subscriptions into their ecosystems, with subscription-based features such as YouTube Memberships, Facebook Subscriptions, and Twitter SuperFollow that let fans pay monthly subscriptions in return for exclusive benefits. However, despite the membership model recently gaining wide importance, it is currently understudied.

Understanding how membership platforms function is becoming increasingly important as they continue to drive this paradigm shift in how creative work is valued and supported. Millions of people already consider themselves full-time creators, and the nature of education, entertainment, and other industries are rapidly changing as creator-fan relationships become more engaged and participatory. Furthermore, the presence of financial transactions stands in stark contrast with the free-to-use nature of most large online platforms, and behavior on membership platforms may be substantially different as a result. To shed light on membership platforms, recent work has applied qualitative methods to detail the experiences of small groups of creators and patrons. For example, a recent study analyzed the campaigns of 40 Patreon creators through the framework of parasocial relational work, in which creators develop one-to-many intimacy with their fans while simultaneously managing social and economic relationships with them (Hair 2021).

However, the structure and dynamics of how membership platforms function on a large scale remain poorly understood. How do patrons distribute their pledges over time? How are creator audiences composed? How have the platforms themselves evolved as the creator economy has rapidly grown? These questions are difficult to answer as they require fine-grained, large-scale, and longitudinal data on how patrons and creators interact. Furthermore, the lack of quantitative studies of membership platforms makes it unclear how to analyze creator-fan dynamics.

The Present Work. In this work, we develop an analysis framework for the study of membership platforms and apply it to the complete set of all hundreds of millions of Patreon pledges from its inception in 2013 until the end of 2020. We analyze Patreon activity from three perspectives: patrons, creators, and the platform as a whole.

We study how Patreon's millions of patrons pledge to creators on the platform. We propose loyalty (continuing to support a given creator), retention (continuing to support any creator), patron tenure (how long the patron has been on the platform), and total monthly pledging amount as four key patron metrics to understand. We find that loyalty and retention are strongly related to the number of pledges a patron makes, but surprisingly, controlling for a patron's number of pledges, total amount pledged on the platform is negatively associated with loyalty and retention. This differs significantly from patterns observed on other platforms, where strength of engagement is essentially universally associated with loyalty and retention. On membership platforms, stronger engagement is heavily influenced by creator differences. On the other hand, pledges made to a creator's higher tiers are more likely to be maintained on average, showing that spending within-creator is a good indicator of loyalty on average.

Examining pledges over time, we find that pledging behavior is fluid and dynamic. As patrons increase in the number of pledges they make, they are more likely to change the set of creators they support over time. A key question for membership platforms is whether patrons engage in budgeting behavior as they decide to pledge to new creators. Since the main consumption resource on membership platforms is money, as opposed to time as on other platforms, it is possible that new pledges compete with existing pledges. We find limited evidence that patrons engage in budgeting, and adding new pledges is associated with increasing total spending, even for highly active patrons.

A defining characteristic of a patron's engagement is how broad their pledging patterns are. Patrons vary between specializing in a narrow topic of interest to pledging more generally to diverse sets of creators. To capture this, we construct a *creator embedding* to capture behavioral similarity between creators, where two creators are similar if they are pledged to by similar patrons. We use this embedding to quantify patron breadth, and find that specialists are more loyal and pledge higher amounts to the creators they support, but churn off the platform more frequently.

We then analyze Patreon activity from the perspective of creators, concentrating on contrasting the audiences of large and small creators by earnings across different content modalities. Interestingly, the highest-earning creators have the largest patron-bases, but the patrons who support them are the least likely to support other creators as well. Large creators also receive higher and more diverse pledge amounts, with audience behaviors varying significantly according to the creator's modality.

Finally, we analyze platform-level metrics that inform how Patreon evolved over time. We observe that activity on Patreon consistently became more diversely distributed across topics over time, as the creator economy itself grew and expanded. During this growth, loyalty and retention remained fairly stable, more creators started to receive earnings, and the platform launched higher-earning creators over time.

Our work contributes the first complete, large-scale analysis of granular, transaction-level data on how patrons and creators connect on membership platforms. We shed light on how patrons manage their pledges, how creator audiences are composed, and how the largest membership platform evolved throughout its history.

Background and Related Work

The creator economy. Social media platforms have democratized media creation and publication, allowing anyone to share their skills, talents, and opinions with the world. This upheaval in how media is distributed represented an opportunity for brands to increase awareness by partnering with popular people seeking to monetize their audiences. But the resulting advertising partnership model had major drawbacks, leading to inauthentic pitches and influencer fatigue. Now, subscription-based platforms such as Patreon, Twitch, and OnlyFans have enabled creators to be directly supported by fans (Forbes 2021). The creator economy has become the fastest-growing type of small business, and more American children want to be a YouTube star (29%) than an astronaut (11%) when they grow up. Membership platforms (MPs) allow for creators to receive monthly earnings from their followers, either in exchange for exclusive content, access to creators, or simply as a means of support (Regner 2020). MPs differ from standard crowdfunding platforms, such as Kickstarter and IndieGogo, in which creators pursue one-time donations to support time-limited projects. On membership platforms, patrons give monthly pledges and receive benefits in return.

Crowdfunding platforms. A rich body of literature has developed to study the dynamics of crowdfunding and related platforms (Barzilay et al. 2018; Mollick 2014; Agrawal, Catalini, and Goldfarb 2014; James et al. 2020). A popular theme in this work is trying to uncover the determinants of success by comparing elements of successful campaigns to those of failed ones (Mollick 2014). Researchers have found that consumers have prosocial motives to help creators reach their funding goals, backers favor more concrete rewards than symbolic ones, and signals about product characteristics help unlock values of investment characteristics (Dai and Zhang 2019; James et al. 2020; Bapna 2017).

Subscription platforms. The literature has also studied the dynamics of subscription platforms such as Spotify and Netflix, and their impact on the media industry (Hogan 2015; Wlömert and Papies 2016; Guardian 2019; CNBC 2018). For instance, subscribing to Spotify cannibalizes consumers' music expenditures (Wlömert and Papies 2016), and subscription platforms have absorbed a majority stake in movie consumption (PWC 2019). Another popular theme in this literature is trying to understand user consumption behaviours and determinants of churn and retention on subscription platforms (Holtz et al. 2020; Anderson et al. 2020; Gomez-Uribe and Hunt 2016).

Membership platforms. Given their recent introduction, membership platforms have not been as thoroughly examined as crowdfunding and subscription platforms, but several qualitative studies have been carried out. One recent paper conducted a thematic analysis of 40 Patreon campaigns to study how digital creatives strike the balance between connecting and transacting with their patrons (Hair 2021). Studies have shown that YouTubers who have an account on Patreon have more frequent video uploads, and that YouTube channels with a greater number of subscribers have greater membership platform revenue generation (Ciechan-kujawa and Gornowicz 2020). The most similar work to our own analyzed transaction-level Patreon data from 2013-2015 to study earning distributions and determinants of successful creators (Regner 2020). Our work considerably extends this analysis by analyzing a much longer transaction dataset and contributes an analysis framework oriented around patrons, creators, and platform-level metrics.

Online consumption. Our work is concerned with the consumption patterns on online platforms, and how this relates to consumer engagement. We draw on previous work on user churn and loyalty (Gomez-Uribe and Hunt 2016; Anderson et al. 2020; Holtz et al. 2020; Hamilton et al. 2017) and how the breadth of a user's preferences impacts their retention and online behaviours (Waller and Anderson 2019; Anderson et al. 2020). This literature has also examined what motivates user participation and consumption on online platforms such as YouTube (Khan 2017; Gomez-Uribe and Hunt 2016) and Spotify (Datta, Knox, and Bronnenberg 2018), and studied platform-level metrics (Burgess and Green 2018). Our research adds to this body of work by analyzing how membership engagement patterns are associated with churn, loyalty, and breadth.

Data

In this paper, we study data from Patreon, a membership platform in which *patrons* can pledge money to *creators* on a monthly basis. Patreon was the first platform to enable fans to directly pay creators, and is now the largest membership platform. We analyze the complete anonymized set of hundreds of millions of patron-creator pledges from the platform's inception in June 2013 until December 2020. Every transaction contains an anonymized patron ID, a creator ID, a pledge amount, and the date (year-month-date) associated with the pledge (i.e. when the creator receives the financial support and the patron receives the benefits associated with the pledge). We also have a dataset of creator metadata containing display name, URL, summary, modality, and a subset of topic categories.

Each creator specifies a set amount of tiers that patrons can choose from, where each tier contains certain benefits and has a monthly donation amount attached to it. For example, the *True Crime Obsessed* podcast, which received payments from tens of thousands of patrons in 2021, offers 3 tiers: $7 gives a patron access to exclusive bonus episodes and pre-sale access to live shows; $9 also gives a patron access to another set of bonus content; and the $13.50 dollar tier offers ad-free content in addition to all of the above.

Activity on Patreon is particularly organic, as the platform does not offer algorithmic or human-curated recommendations, and there is no notion of a patron-patron social network. As such, patrons can only arrive at a creator's page via search or via an off-platform link. This lack of algorithmic confounding allows us to study behavior driven by more organic user choices.

We study Patreon from three perspectives: patron engagement, creator audiences, and platform-level metrics.

Patron Engagement

Membership platforms are driven by fans engaging with creators through financial pledges, and we begin by analyzing this behavior. On the vast majority of online platforms that have been studied to date, engaging with creators has no monetary commitment associated with it. Since patron engagement on Patreon involves financial support, the basic mechanisms governing patron engagement could be quite different from established patterns of activity on platforms such as YouTube. Here we examine how patrons choose to allocate their support.

Loyalty and retention. How do patrons change their engagement over time? Here we study *loyalty*, patrons continuing to support creators they have pledged to, and *retention*, patrons continuing to actively pledge on Patreon.

Given the set of all pledges R^t in month t, consider patron p and let $r^t \in R^t$ denote the set of pledges this patron makes in month t. The i^{th} pledge r_i^t is a tuple (C_i^t, D_i^t) with C_i^t denoting the creator receiving the pledge and D_i^t denoting the amount pledged. Then we say p was *loyal* to C_i^t if p is still pledging to C_i^t at time $t+k$ (i.e. if $C_i^t \in C^{t+k}$). Given all the pledges on the platform, the loyalty rate on the platform at t is thus the fraction $|R^t \cap R^{t+k}|/|R^t|$ of pledges that are maintained at $t+k$. Furthermore, we say patron p was *retained* if $C^{t+k} \neq \emptyset$ (i.e. they did not churn). The platform's retention rate at t is thus the fraction of patrons with $C^{t+k} \neq \emptyset$. Throughout this paper we report relative engagement rates (percentage point difference from mean rate). We also set $k = 3$ months, but other choices yield qualitatively similar results.

How do patron loyalty and retention vary with activity ($|C^t|$, the number of creators pledged to) and amount spent on pledges ($\sum_i D_i^t$)? To allow comparisons across different activity levels, we group patrons into dollar spending deciles based on their activity level that month—those who spend the least given $|C^t|$ are in decile 1, and those who spend the most given $|C^t|$ are in decile 10. For example, among patrons who pledge to one creator, those who spend $1 are in the 1st decile and those who spend over $20 are in the 10th decile; while among patrons who pledge to 5 creators, those who spend a total of $2-10 are in the 1st decile and those who spend over $58 are in the 10th decile.

We find that the number of creators a patron supports ($|C^t|$) is a clear signal of how engaged they are. This measure of activity is directly proportional to patron retention (Figure 1a) and loyalty to creators (Figure 2a). Patrons who support many creators are more likely to stay on the platform and to continue supporting their creators, while uni-patrons are the least likely to stay on Patreon or be loyal to their creators. The effect size is particularly large for retention; uni-patrons are retained 6 percentage points less than average, while patrons who pledge to 10 campaigns are retained around 14.7 percentage points more than average, a striking

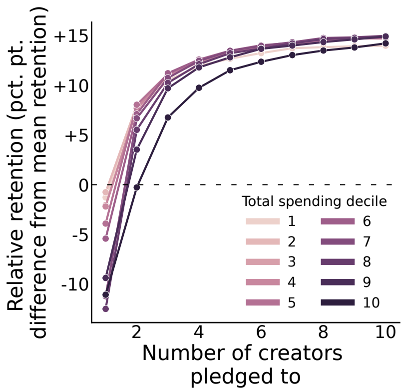

Figure 1: Relative retention (pct. pt. difference from mean retention) per patron's spending decile given the number of campaigns they pledge to (1 is the lowest, 10 is the highest).

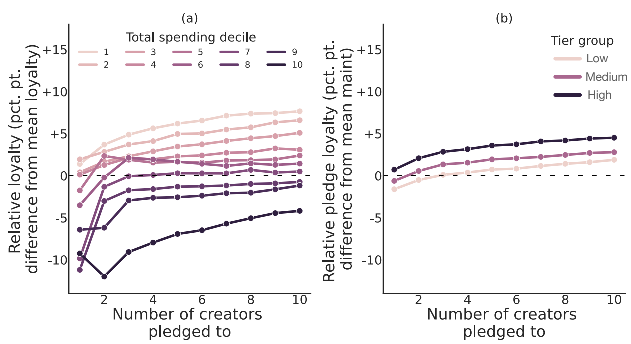

Figure 2: (a) Relative 3-month loyalty (pct. pt. difference from mean loyalty) per spending decile and activity. (b) Relative 3-month pledge maintenance (pct. pt. difference from mean maintenance) per patron's pledging tier and activity.

20.7 pct pt difference. It is less pronounced but still significant for loyalty; uni-patrons continue pledging to a creator 3.8 percentage points less than average, while patrons who pledge to 10 campaigns maintain their pledges 2.1 percentage points more than average (a 5.9 pct point difference).

Interestingly, controlling for the number of campaigns pledged to, higher spenders on the platform are more likely to churn (Figure 1) and are less loyal (Figure 2a). This effect is more pronounced for loyalty than it is for churn. The difference in churn rates between the highest and lowest spenders among those who pledge to one campaign is 10 percentage points, but this difference is only 1–2 percentage points for those who pledge to 10 campaigns. The effect is consistent for loyalty: the highest spenders across creators are consistently around 10 percentage points less loyal than the lowest spenders, no matter whether they pledge to 1 or 10 campaigns.

This phenomenon stands in stark contrast with behavior on other online platforms, where stronger forms of engagement are generally associated with higher loyalty and lower churn (Hamilton et al. 2017). Why are the seemingly most engaged patrons also the least loyal? To resolve this puzzle, we also measure how pledges *to the same creator* vary across pledging tiers. For each creator, we group their patrons into 3 equal-sized tier groups (low, medium, and high) based on the amount each patron pays that creator. This analysis shows that higher spenders within-creator are more likely to be loyal on average, with a 3 percentage point difference between the highest pledges and the lowest (Figure 2b). Thus the original pattern was driven by differences *between* creators: creators with larger expected pledges also inspire lower loyalty rates.

Loyalty over time. Loyalty is a fundamental property of the fan-creator relationship. Since patrons vary in how long they've been on the platform, and how long they've been pledging to a creator, we study how a patron's tenure (time on the platform) and their pledge's tenure (time pledging to the same creator) impact engagement.

Consider a pledge between a patron p and a creator c. At time t, we define the *pledge tenure* as the difference between t and the date of p's first pledge to c. Similarly, a patron's *platform tenure* is the total calendar time since their first pledge on Patreon. We measure patron loyalty rate as a function of platform tenure and report the rate's percentage point difference from the mean. As can be seen in Figure 3a (orange line), a patron's platform tenure is directly proportional to their likelihood of maintaining their pledges—platform veterans tend to be more loyal. To measure loyalty as a function of pledge tenure, we compute the fraction of pledges within each tenure-group that are maintained at time $t + k$ and again report the percentage point difference from the mean. Figure 3a (purple line) shows that a pledge's tenure is also directly proportional to its likelihood of being maintained, and is more indicative of loyalty for later months—longstanding fans are most likely to remain loyal. The effect size is large for both measures, with more than a 24 percentage point difference between the newest (tenure = 1 month) and oldest patrons (tenure = 24 months).

To understand how loyalty falls off with time, we measure the likelihood of a pledge lasting for at least k months. As can be seen in Figure 3b, short loyalty streaks are far more common than long ones, with a 33.5 percentage point drop in expected loyalty between a patron being loyal for 1 month compared to 12. The effect is also stronger for less active patrons. The difference across activity becomes more prominent for larger values of k; patrons who support ten campaigns are 6.7 percentage points more loyal than those who support only one.

Patron Budgeting

A key distinguishing feature of membership platforms is that they involve financial transactions. Subscribing to a new creator on traditional online platforms is low-cost—for example, there is no technical limit to the number of users one can follow on Twitter or YouTube. But on membership platforms, patrons need to be more discerning in who they support. Do patrons manage their portfolio of pledges to stay within a given budget, and does the addition of new pledges lead to *budgeting behavior* in which patrons drop or swap existing pledges?

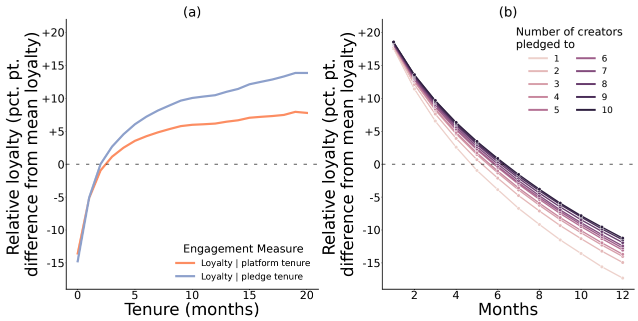

Figure 3: (a) Relative 3-month loyalty (pct. pt. difference from mean loyalty) given the number of months the patron has been on the platform (orange) and the months they've been pledging to that creator (purple) and whose start year $\geq$ 2019 . (b) Relative k-month loyalty (pct. pt. difference from mean loyalty).

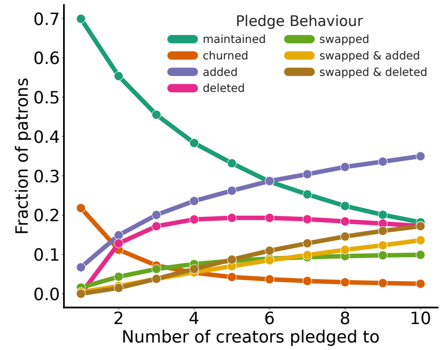

Figure 4: How patrons changed their pledging activity (either maintaining, churning, adding, deleting, swapping, swapping and adding, or swapping and deleting) over the course of three months, averaged over 2019–2020.

The answer to these questions have deep implications for membership platforms. If patrons are consistent in their pledging behavior, recommendation systems might be a good strategy to encourage more exposure to creators and increase activity on the platform. But if patrons exhibit budgeting behavior, this is less clear. We investigate the degree to which patrons budget to help creators manage their engagement strategies and inform membership platforms on the potential impacts of encouraging more activity through recommendations.

Pledging strategies. To answer these questions, we conduct an in-depth analysis of how patrons change their sets of pledges over time. For every patron, we compare their set C^t of creators pledged to at time t with their set C^{t+3} of creators pledged to three months later, and classify their behavior based on the differences between these sets. We categorize pledge changes as follows: *maintenance* is when a patron makes no change to their pledges ($C^t = C^{t+k}$); *churning* is when a patron leaves the platform ($C^{t+k} = \emptyset$); *adds* occur when a patron starts pledging to creators they don't currently support ($C^{t+k} \supset C^t$); *deletes* occur when patrons stop pledging to a creator but remains on the platform ($C^{t+k} \subset C^t$ and $C^{t+k} \neq \emptyset$); and *swaps* are when a patron replaces a pledge to one creator with a pledge to another ($C^{t+k} \not\supset C^t$ and $C^{t+k} \not\subset C^t$). Note that swaps may occur along with adds or deletes, depending on whether $|C^{t+k}| > |C^t|$ (adds) or $|C^{t+k}| < |C^t|$ (deletes).

We apply this taxonomy of temporal pledging differences to all patrons active in 2019–2020. We report on the fraction of patrons at t that fall into each group (as the groups are mutually exclusive), comparing their pledging behavior three months apart (Figure 4). Several patterns emerge from this analysis. First, as $|C^t|$ grows, who patrons support is more likely to change—the maintenance curve decreases monotonically. This is intuitive, as the likelihood that a patron decides to keep their exact set of pledges should decrease with the complexity of their pledge set. However, the amount of change is surprising. Only patrons who make 1 or 2 pledges are more likely to maintain their exact set than to change it. Second, we see limited evidence of budgeting behavior. Strict addition (purple) monotonically increases while strict deletion (magenta) and churning (red) eventually decrease. Swaps that include deletions (brown) outpace swaps that include adds (yellow), but they are consistently close in frequency. Finally, complex changes—those that include swaps, meaning pledging sets at time t and $t+3$ are incomparable—become more prominent as $|C^t|$ increases.

Budgeting strategies: Platform perspective. When a patron pledges to a new creator, what happens to their overall activity on the platform? To more directly investigate this budgeting question, we analyze all patrons who made at least one new pledge between t and $t+k$ and measure the average change in their number of pledges ($|C^{t+k}| - |C^t|$) and total spending ($\sum_i D_i^{t+k} - \sum_i D_i^t$). We analyze $k = 3$ in this analysis to allow time for budgeting to take place, although other choices of k yield qualitatively similar results. As shown in Figure 5, the average change in the number of campaigns supported and total money pledged goes down with $|C^t|$—the more pledges a patron starts out with, the more likely they are to exhibit budgeting behavior when they add a new pledge. However, the net change is always positive, even for patrons who have 10 pledges at time t.

Budgeting strategies: Creator perspective. Budgeting is thus limited from the platform perspective, but what about from the creator's point of view? When a patron adds a new pledge, what happens to the creators they were already pledging to—and how does this compare to when this patron *doesn't* add a new pledge?

For every patron who made a new pledge between times t and $t+k$, we focus on how this new pledge impacted that patron's existing pledges (i.e. C^t). We report both the probability of each existing pledge being dropped at $t+k$ ($\frac{|C^t \cap C^{t+k}| - |C^t|}{|C^t|}$) and the mean change in pledge amount across all C^t (D_i^{t+k} minus D_i^t). We call this the *adds* group. However, simply looking at these changes for the *adds* group in isolation conflates two reasons why patrons may drop existing pledges: they could be budgeting (i.e. drop-

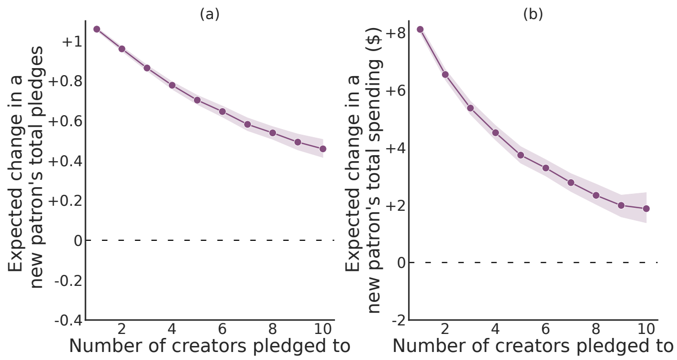

Figure 5: (a) Average change in a patron's number of creators pledged to one year later given that they pledged to at least one new creator during the year. (b) Average change in patron spending one year later given that they pledged to at least one new creator during the year.

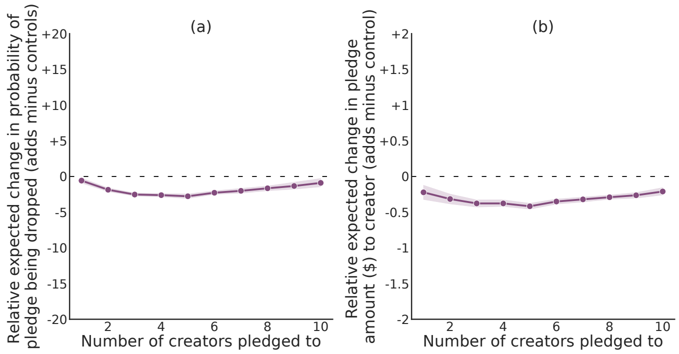

Figure 6: Difference between the adds and control groups' (a) probability of creator's pledge being dropped and (b) expected change in pledge amount, for 2019–2020 data.

ping an existing pledge because they added a new pledge), or they could be dropping the pledge for any of the other reasons why pledges don't necessarily last forever (i.e. they churned, they only intended to pledge once, etc.). So whatever changes we observe in the adds group cannot just be attributed to budgeting.

To control for this, we compare the changes we observe for the *adds* group with the changes we observe in a *control* group of patrons who didn't make a new pledge by time $t+k$. We plot the difference between these two groups in Figure 6. Between 2019–2020, we find that budgeting is very limited: a creator whose patron adds a new pledge only has a 1.8 percentage point higher chance of having their pledge dropped (and expects to lose an average of \$0.315 more) than if their patron didn't make a new pledge Figure 6. The difference between the *adds* group and *control* group is extremely close to zero across all activity levels, revealing that creators whose patrons add a new pledge at $t + k$ will not differ too much from those whose patrons don't make a new pledge.

Patron Breadth

Patrons vary in the breadth of their pledges, ranging between pledging to a narrow group of related creators and pledging more generally to a group of more diverse creators. A patron's breadth of engagement is a reflection of their interests and how they use and derive value from the platform. In other platforms, activity diversity has been related to user loyalty, engagement, and retention (Waller and Anderson 2019; Anderson et al. 2020). But how is patron breadth related to key user metrics on membership platforms, where loyalty and longevity come at a financial commitment? To answer this question, we measure how a patron's spending, platform retention, and creator loyalty are related to their breadth of support on Patreon.

Quantifying patron breadth. Measuring the breadth of a patron's pledging activity is a difficult problem, as it requires a fine-grained, scalable, and consistent measure of creator similarity—patrons who pledge to "similar" creators pledge narrowly, whereas patrons who pledge to "dissimilar" creators broadly. But how do we measure creator similarity? To solve this problem, we adopt a community embedding approach, where creators are embedded in a high-dimensional space based on how patrons pledge to them (Waller and Anderson 2019). We adapt the word2vec word embedding algorithm to patron-creator pledging data to generate community embeddings. Treating creators as "words" and patrons who pledge to them as "contexts", we embed communities into a high-dimensional vector space. Similar to (Waller and Anderson 2019), we use the skip-gram model with negative sampling and train over all pairs (c_i, p_j) of patrons p_j pledging to creator c_i . Creators are then close in the space if many patrons pledge to them both. For any pair of creators, we calculate their similarity with the cosine similarity of their respective vector representations. We visualize a two-dimensional t-SNE projection of a portion of the resulting creator embedding in Figure 7, where every point represents a creator and is colored by topic. Even though topics were assigned by a separate independent, content-based manual procedure, the purely behavioral, pledge-based embedding clusters into the topics strikingly well.

The creator embedding provides us a method of measuring behavioral similarity between every pair of creators, but we still require a method of calculating the breadth of an arbitrary set of creators that a patron pledges to. We use the *generalist-specialist score* (GS-score), which is simply the average cosine similarity between all of a patron's creators (Waller and Anderson 2019), and has been used to capture behavioral similarity on Reddit, Spotify, and GitHub (Waller and Anderson 2019; Anderson et al. 2020). This captures the intuition that *specialists* are those who pledge to highly similar creators that are clustered close together in the space (high GS-score), whereas *generalists* are those who pledge to dissimilar creators that are far apart in the space (low GS-score). To illustrate, a typical generalist would pledge to creators that are far apart in the space—e.g. in Figure 7 an ASMR creator, a Lithuanian creator, and a Minecraft mod creator. In contrast, a typical specialist would pledge to creators who are very close in the space (although

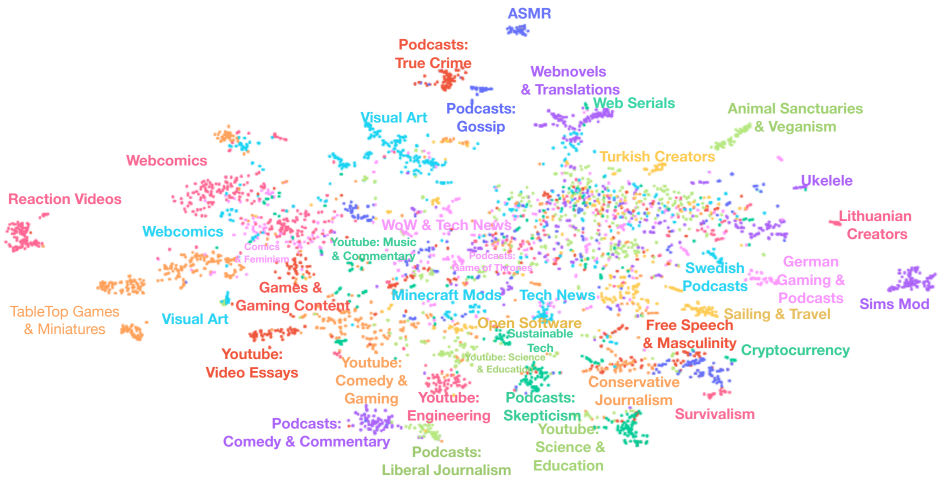

Figure 7: Two dimensional t-SNE projection of a subset of safe-for-work (SFW) creators, labelled by topic.

not necessarily classified in the same topic). For example, a YouTuber focusing on mountain bikes, a group devoted to creating public bike parks, and a podcast about bike culture.

The creator embedding is also a good illustration of the breadth of topics present on the platform. The subset of topics shown in Figure 7 give us an inside look into the range of content being shared on the platform. Creators share everything from Sustainable Tech and Cryptocurrency to Visual Art and Comics Feminism, thus revealing the wide range of content that Membership platforms support and creator economy contains.

Patron breadth, loyalty, and retention. The distribution of GS-scores reveals that patrons exhibit a wide range of pledging styles, with a tendency towards specializing around a particular area of interest. Multi-patron GS-scores have a large tail towards the maximum GS-score of 1, with the majority of users scoring above 0.8. In comparison to a null hypothesis where patron-creator edges are shuffled, Patreon has a more specialised user base. This is intuitive, as multi-patrons tend to pledge to more similar creators than if they had drawn them at random.

Previous applications of this GS-score methodology on Reddit, GitHub, and Spotify have shown that specialists are more loyal to their communities, while generalists are more likely to stay on the platform (Waller and Anderson 2019; Anderson et al. 2020). Is this the case on a membership platform like Patreon?

In Figure 8a, we plot retention (relative to the global baseline retention rate) as a function of number of pledges and GS-score, and in Figure 8b we plot loyalty (relative to the global baseline loyalty rate) as a function of number of pledges and GS-score. Consistent with previous work, more generalist patrons are more likely to be active on the platform 12 months later (Figure 8a), but are less likely to keep pledging to the same creators (Figure 8b). The most extreme specialists, however, are not impacted by engagement level.

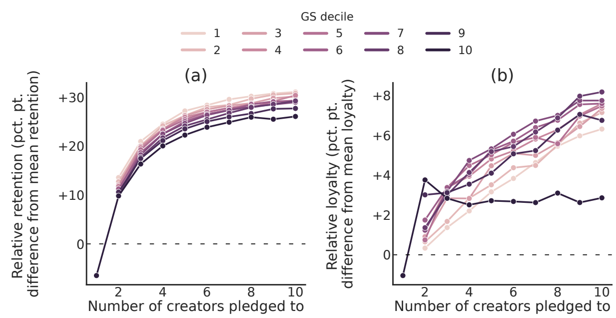

Figure 8: (a) Patron 12-month relative retention (pct. pt. difference from mean retention) as a function of number of pledges and GS-score. (b) Patron 12-month relative loyalty (pct. pt. difference from mean loyalty) as a function of number of pledges and GS-score.

Controlling for the number of pledges, specialist users are more likely to have a strong interest in a particular area, which explains their narrow pledging and high loyalty. But as their narrow behavior could also indicate that they have only discovered few use cases for them on the platform, and that this use case could potentially be satisfied elsewhere. Hence, they also churn more often from the site. On the other hand, generalist users are more likely to be satisfying more of their interests on the platform, and are correspondingly less likely to churn from the site. However, their breadth of interest is associated with lower pledge amount per creator and less loyalty to existing pledges.

Creator Audiences

We now turn to understanding Patreon from the creators' perspective. Success as a creator on a membership platform can spell the difference between being able to support a creative career or not. Hence, it is important to understand the

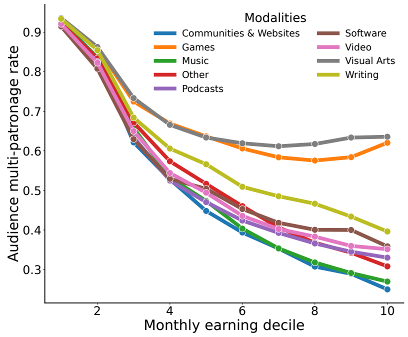

Figure 9: Fraction of a creator's audience who support ≥ 2 creators per creator earning decile at t

structure and dynamics of membership platforms from the creator's point of view. Here we examine characteristics of creators' audiences and what distinguishes large from small creators.

First we establish some notation. Consider a creator c at time t and T^t denote the set of pledging tiers patrons can choose from. Let C^t denote the set of pledges made to c in month t, where the i^{th} pledge is a tuple of (P_i^t, D_i^t) with P_i^t denoting the patron giving the pledge and $D_i^t \in T^t$ denoting the tier amount pledged. Then $m_t = \sum_i D_i^t$ is c's total earning in month t and $s_t = |C^t|$ is c's number of patrons pledging in month t.

Creator audience size. To understand what differentiates the audiences of larger and smaller creators, we examine how they grow and shrink over time. Throughout this section, we group creators into deciles based on their total monthly earning $\sum_i D_i^t$, with 1 being the lowest and 10 being the highest.Thus, a creator's size refers to their monthly earnings such that a larger creator is in a higher earning decile. We also group creators according to the *modality* of their content. Modalities falls into one of these categories: {*Video, Podcasts, Music, Games, Writing, Visual Arts, Software, Communities & Websites,* and *Other*}. Modalities are determined by a machine learning model that classifies the type of a creator's content based on creator features and natural language processing of the text in various parts of the creator page.

To compare how creators' audiences change over time, we consider each creator's sets of patrons P^t and P^{t+k} at times t and $t+k$, and measure a creator's *maintenance* as $\frac{|P^t \cap P^{t+k}|}{|P^t|}$, *churned* patrons as $\frac{|P^t - P^{t+k}|}{|P^t|}$, and *change in size* as $\frac{|P^{t+k}| - |P^t|}{|P^t|}$.

We find that at any one time, all types of creators get the majority of their pledges from maintained patrons—loyal patrons who have paid them in the past. All creator tiers exhibit positive growth rates in audience size, revealing that all tiers are experiencing growth on the platform regardless of churn. Interestingly, smaller creators exhibit higher relative growth in audience size, revealing that it is also possible for newcomers on the platform to gain prominence.

Creator audience engagement. What type of audiences do creators attract? Here we study creator audiences through *multi-patronage*, the fraction of patrons at time t that are also pledging to at least one other creator ($|\{P_i^t \mid i$ s.t. $|C_i^t| \geq 2\}| \;/\; |P^t|$), and *modality*, the modality through which the creator produces (videos, music, writing etc..).

We find that larger creators attract audiences with lower rates of multi-patronage (Figure 9). Surprisingly, the audiences of the largest creators pledge to other creators 40% of the time, while the audiences of the smallest creators are multi-patrons 90% of the time —over a 50 percentage point difference between the largest and smallest creators. This reveals that large creators often inspire "one-and-done" patrons, those who come to the platform only for them. We also find that the type of content is correlated with multi-patronage rates: Visual Arts and Games have the highest multi-patronage rates, while Music and Communities and Websites have the least. The differences between modalities increase with creator size, thus revealing that a creator's modality may relate to several creator outcomes.

Creator Earnings

Membership platforms are unique in that they revolve around direct financial payments between creators and patrons. In this section, we characterize how support for a creator is distributed among their patrons.

Mean pledge amount. Larger creators generate more earning on the platform, but this could be due to bigger pledges from their fans, or simply because they command larger audiences. In Figure 10a, we compute mean pledge amounts for every creator and examine how this varies with creator size and modality. There is a strong correlation between creator earning decile and mean pledge amount, with higher-earning creators receiving higher pledge amounts on average. Although the direction of this relationship is expected, the magnitude of the effect is striking. The highest-earning creators receive almost \$15 per pledge, whereas the lowest earners receive only \$2 per pledge. Thus not only do high-earning creators attract larger audiences, they also attract far larger pledges per patron. We also see that the difference in pledge amounts across modalities becomes more prominent with larger creators, with more than a \$12 difference in mean pledge amounts between modalities.

Tier structures. Patreon's pledging model allows creators to specify as many pledging tiers as they like, and patrons can pledge to any of the available tiers. Each tier identifies a set of benefits that patrons receive upon pledging a certain amount. How are creator earnings distributed among their patrons? On one extreme, creators could receive equal pledges from all of their patrons, and on the other extreme they could receive all of their payments from a single patron.

We calculate the Gini coefficient of a creator's pledges, as shown in Figure 10b. First, we observe that the mean Gini coefficient is at most 0.45, indicating a moderate but not extreme amount of inequality. Interestingly, there is a clear relationship between creator earning decile and their patrons' pledging distributions—higher-earning creators receive more variable pledge amounts than low-earning cre-

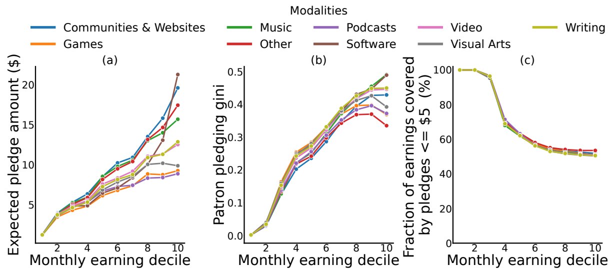

Figure 10: (a) Expected pledge amount per creator earnings decile. (b) Gini coefficient of per-creator pledge distribution. (c) Fraction of earnings covered by small ($\leq$ \$5) pledges as a function of creator earnings decile.

ators. The higher, more variable pledge amounts that these creators attract form an interesting contrast with their lower patron loyalty rates. This trade-off may occur because patrons want access to time-sensitive benefits from higher tiers, or because they prefer to pay their favorite creators with higher pledge amounts for shorter periods of time. Finally, although small creators typically specify several tiers that patrons can pledge to, their pledges are very evenly distributed, resulting in average Gini coefficients near 0. This indicates that virtually all of their patrons choose to pledge from the same pledging tier. The Gini coefficients do seem to differ across modalities, but the differences are not as pronounced as that of expected pledge amounts or multi-patronage.

Pledge sizes. Given the low earnings of the smallest creators, the single tier that their patrons pledge to is presumably a low one. To check this directly, we measure the fraction of creator earnings that are accounted for by "small" pledges, which we operationalize as $\leq$ \$5. Figure 10c shows that lower-earning creators indeed receive virtually all of their earnings from small pledges. Between deciles 3 and 4, there is a stark drop down to 70% earnings covered by small pledges—indicating that the 4th decile is where creators start attracting more substantial pledges. This percentage smoothly drops to around 50% for the highest-earning creators, who derive their earnings from an equal split of big and small pledges. This phenomenon seems to be true for all creator modalities, with little difference in behaviour across modality groups.

Platform-Level Metrics

The analyses we have conducted thus far studied how Patreon operates from the perspective of patrons and creators. Here, we take a system-level view of Patreon as a platform. How has the largest membership platform grown and evolved over time, and how has the creator economy ecosystem within it changed?

Evolution of retention and loyalty. Given that Patreon pioneered creator-fan memberships and remained the largest membership platform through the end of our data time window, we can shed light on how the dynamics of the creator economy evolved over time. How have engagement mea-

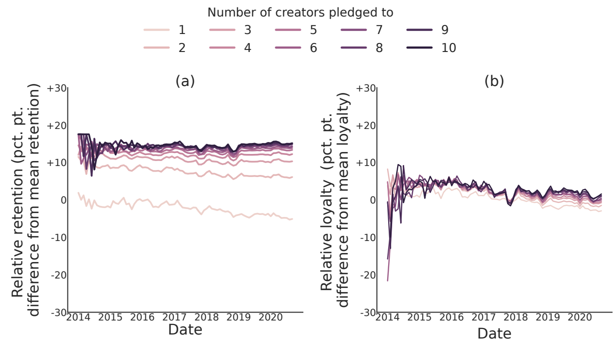

Figure 11: (a) Relative retention (pct. pt. difference from mean retention) over time and across number of creators pledged to (b) Relative loyalty (pct. pt. difference from mean loyalty) over time and across patron activity levels

sures changed over time? We measure loyalty and retention (relative to the global average) across activity levels and time. Figure 11 shows that in mid-2014, the platform's engagement began to stabilize. Users were maintaining their pledges more consistently, with loyalty slightly decreasing over time. A strong relationship emerged between activity level and retention and has been in place ever since: the more creators a patron pledges to, the more likely they have been to maintain those pledges. The engagement dynamics of the creator economy have been relatively stable, even during times of explosive growth.

Distribution of earnings. How are earnings distributed among creators, and how has this distribution changed over time? We measure earning distribution by calculating how the total earnings from the platform's inception has evolved across its original deciles. We calculate earning deciles at the start of the platform's history — which monthly earning amounts would have qualified a creator to be in a particular earning decile in December 2013. Then we measure the total amount generated by each earning decile (all creators who make that particular band of amounts) over time.

Let $I(t)$ be the set of earning intervals for each decile i at time t. Then each earning decile i at time t is composed of creators with monthly earnings $m_t \in (min_{(i,t)}, max_{(i,t)}]$. We calculate $I(2013)$ and classify all future creators based on which decile they would have been in 2013. Thus, a creator c_t with monthly earnings m_t will be in the decile i where

$$min_{(i,2013)} \leq m_t < max_{(i,2013)}$$

We plot the total group earnings over time in Figure 12. We find that every part of the platform's earning groups have grown substantially since 2013, with each decile generating more earnings, revealing that the creator economy has been allowing more creators to receive earnings and launching higher-earning creators over time. We also find that the amount of money generated by these groups is very substantial, revealing the prominence and importance of membership platforms and the creator economy.

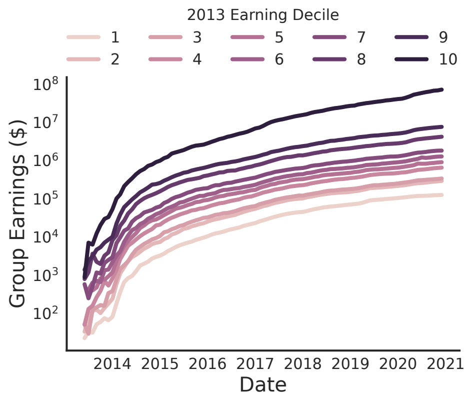

Figure 12: Group earnings generated by creators who fall into each 2013 earning decile

Modality	Relative Loyalty (pct. pt)	Relative Pledge ($)	Relative patron activity
Podcasts	+6.7	−0.5	−1.7
Writing	+3.9	+0.3	−0.3
Music	+1.9	+1.4	−1.4
Communities & Websites	+0.3	+2.2	−1.8
Video	+0.1	+0.2	−0.8
Software	−1.5	+0.2	−0.9
Visual Arts	−2.6	−0.5	+2.7
Other	−3.2	+1.5	−1.0
Games	−9.5	−0.7	+2.0

Table 1: Relative loyalty (percentage point difference compared to the mean loyalty rate), expected pledge ($ difference from mean expected pledge), and relative patron activity (number of creators pledged to, diff from mean activity) for creators who fall into each modality, sorted in descending order by relative loyalty.

Engagement by content type. As the largest membership platform, Patreon supports a wide variety of creators. To dive deeper into this, we analyze the type of content supported on the platform and the differing patron behavior across these modalities. In Table 1, we show how patrons engage with creators in the major modalities. The loyalty creators inspire is heavily related to their modality, with more than a 13.2 percentage point difference between the modality with the most loyalty (Podcasts) and that with the least (Games). Modalities differ in their expected pledge amounts, as different types of creators differ in the types of benefits given to patrons. The *Communities & Websites* modality has the highest expected pledge amount, which may be reflective of pledging that is more transactional in nature. Furthermore, patrons' activity levels differ greatly across topics. *Gaming* and *Visual Arts* boast patrons who follow 2 or more creators above average.

These results reveal the importance of taking into account a creator's topic when looking at engagement measures on the platform. Learning creator topics and modalities seems to be an important aspect for understanding platform behaviour, and platform designers should invest more time in classifying different creator modalities. Moreover, these results remind creators how much impact the type of content they produce has on their performance in the creator economy. A worthy extension to this analysis would be to understand how these differences across topics relate to behaviours on other social media platforms.

Topic cohesion. Our final platform-level measure takes a closer look at how earnings are distributed across various topics. Patreon originally formed around a band and initially gained traction with YouTubers, but then branched into the creator economy more generally. How did the evolution of earnings reflect this topical growth?

We study platform cohesion by examining how revenue is distributed across different topics on the platform. Each creator is associated with a Patreon-specified topic (e.g. YouTube, Podcast, Visual Arts). We calculate the Gini coefficient of earnings over topics to understand how thematically diverse pledging on Patreon is, and show how it evolved over time in Figure 13a. The Gini coefficient of a topic's revenue is a standard measure of distributional inequality. Pure equality results in a Gini coefficient of 0, pure inequality results in a Gini coefficient of 1, and in general the index lies between these two extremes. In the platform's first 2 years, the revenue across topics was becoming more concentrated, creators were discovering Patreon for the first time and most of them were from the YouTube world. However, ever since mid-2015, there has been a steady, minor decrease in the platform topic Gini coefficient from 0.80 to 0.72, indicating that aggregate pledging on Patreon has supported a slightly more diverse set of topics over time.

However, although the earnings distribution has become slightly more distributed across topics, it has become significantly more concentrated *within* topics over time. The Gini coefficient across creators within a topic has steadily and substantially increased from 0.30 in 2015 to 0.60 in December 2020 (Figure 13b). As Patreon has grown to encompass all manner of creative industries, the top creators within each domain have accounted for a larger percentage of its earnings such that the platform earnings are more heavy-tailed.

Discussion

In this work, we presented an analysis framework for the study of membership platforms and applied it to a dataset comprising all Patreon pledges from 2013–2020. The research presented aims to shed light on the structure and dynamics of how membership platforms function, which has not yet been explored in a large-scale, fine-grained manner.

Patron engagement. Our first contribution is an understanding of patron engagement. We find that the ways in which people engage with creators on Patreon systematically differ from other platforms, but higher pledge tiers within-creator are more likely to be maintained on average. For example, controlling for number of pledges, patrons who pledge

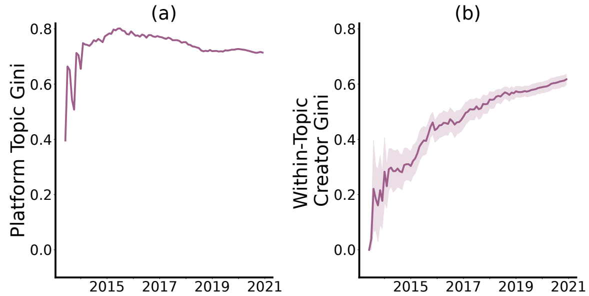

Figure 13: (a) Gini coefficient of platform earnings over topics, over time. (b) Gini coefficient of within-topic earnings, averaged over topics, over time.

higher amounts across creators are less likely to be loyal and less likely to remain on the platform. Since engagement on membership platforms involves a financial commitment, stronger forms of engagement are more difficult to maintain.

Moreover, since membership platforms are still in their infancy, they have the unique challenge of most users being in low-data "cold-start" regimes. This makes it difficult to develop accurate models of users' preferences. Based on these results, membership platforms could develop alternative methods of learning user preferences, such as leveraging viewership data or off-platform audiences. High-value initial pledges could counter-intuitively be offered easier paths towards lower-value but more sustainable long-term pledging amounts.

Since pledge activity is most predictive of continued support, developing ways for patrons to discover new creators might greatly help convert uni-patrons to multi-patrons. Although higher engagement is associated with more loyalty and platform retention, higher engagement requires an added financial commitment. We find that more engaged patrons are more likely to employ budgeting strategies on the platform but their net changes in activity and spending are typically positive. We also find that creators whose patrons add a new pledge aren't impacted too differently from those whose patrons don't make a new pledge, thus showing that the creator economy cultivates an environment where users are able to increase engagement while keeping their previous pledges largely intact. These results also suggest that introducing recommendations to membership platforms could prove beneficial. Patreon activity is currently entirely organic, but recommendations could help introduce patrons to other creators they would want to support. According to our analysis, the new pledges these recommendations might cause would have only limited effects on existing pledges, although this should be further validated experimentally.

Our study of patron breadth showed that using a creator embedding and quantifying breadth with the GS-score is a strong methodology for capturing how patrons use and derive value from the platform. Pledging to a broad set of creators correlates with remaining on Patreon and dropping existing pledges more often. This highlights a tension at the heart of membership platforms: creators are incentivized to attract specialists, whereas the platforms are incentivized to attract generalists. These incentives will be important to take into account when designing recommendation systems since recommending more specialized data might promote more loyalty, but recommending more general content could encourage more retention.

Creator audiences. Our examination of creator audiences showed that creators' maintenance rates are far greater than their churn rates across all creator sizes. This suggests that creators on membership platforms cultivate a value-for-value exchange that satisfies their supporters' needs. However, the difference in audience growth and audience engagement across creator sizes suggest that the audiences of large and small creators significantly differ in their composition and behavior. Large creators, those that typically have large fan-bases outside of Patreon, attract many patrons, but these supporters are less likely to pledge to other creators. The different earning structure between creator sizes reveals the difference in value that larger creators are able to inspire.

Platform-level metrics. Finally, we have contributed several platform-level metrics to understand the growth of the entire ecosystem. Patreon rapidly grew and expanded in breadth of creator topics over time, filling many niches in the creator economy as creators found audiences in all manner of industries. As it grew, more creators started entering the creator economy and generating earnings. Loyalty and retention remained consistently high once the platform was adopted, revealing that membership platforms can support the creator economy in a sustainable manner. Creators have grown to produce content through many different modalities, and these modalities play a role in each creator's expected engagement rates, earnings, and activity. Finally, the platform has become more diverse over time, more revenue has been distributed across different topics with more concentration within-topics for specific creators.

Implications. Since this emergent type of platform has never been studied on a systemic level, describing the structure and dynamics of how Patreon functions within our three-part analysis framework will give researchers, creators, users, and the general public a detailed inside look at how this important new type of platform operates. This work will help provide insight to membership platform designers into how to prolong tenure on the platform and improve loyalty, and suggests several important implications such as budgeting behavior being limited. This work will also help creators gain a better understanding of user trends and behaviors. Finally, this paper will give computational social scientists and economists a deeper understanding of this emergent type of market.

The metrics and analysis techniques introduced give researchers a framework for studying membership platforms in the future. This work has revealed several covariates that impact important metrics on membership platforms. For example, we now understand how dollar spending, activity level, tenure, and breadth relate to loyalty and retention rates. We also learned that a creator's size and topic are significantly associated with their audience's ecosystem.

Limitations and future work. Our work is limited by the absence of data regarding the types of benefits creators give patrons in return for their financial commitments. Further, although Patreon has no recommendation system, discovery on this platform may be heavily dependent on off-platform recommendations. Finally, the results of our work may vary from creator to creator, since we're reporting on trends in aggregate. Future directions of this work could include understanding the impact of recommendation systems on membership platforms, understanding how a patron's engagement relates to their off-platform behaviour (e.g. in relation to YouTube subscriptions), exploring how the COVID-19 pandemic impacted the creator economy, and developing engagement prediction algorithms based on the specified analysis framework.

Our work introduces an analysis framework for the study of membership platforms and shows how they form a viable online infrastructure for the rapidly growing creator economy. Our findings reveal how patrons decide to pledge to creators, how creator audiences are composed, and how the platform itself has evolved. We hope that these metrics and results inform the design of future membership platforms, and lead to even more connections between the creative class and their supporters.

References

Agrawal, A.; Catalini, C.; and Goldfarb, A. 2014. Some Simple Economics of Crowdfunding. *Innovation Policy and the Economy* 14: 63–97.

Anderson, A.; Maystre, L.; Anderson, I.; Mehrotra, R.; and Lalmas, M. 2020. Algorithmic Effects on the Diversity of Consumption on Spotify. In *Proceedings of The Web Conference 2020*, 2155–2165. ACM.

Bapna, S. 2017. Complementarity of Signals in Early-Stage Equity Investment Decisions: Evidence from a Randomized Field Experiment. *Management Science* 65(2): 933–952.

Barzilay, O.; Geva, H.; Goldstein, A.; and Oestreicher-Singer, G. 2018. Equal Opportunity for All? The Long Tail of Crowdfunding: Evidence From Kickstarter. *SSRN Electronic Journal* ISSN 1556-5068.

Burgess, J.; and Green, J. 2018. *YouTube: Online Video and Participatory Culture*. John Wiley & Sons.

Ciechan-kujawa, M.; and Gornowicz, J. 2020. Success determinants of sports projects financed with donation crowdfunding. *Journal of Physical Education and Sport* 2020(5).

CNBC. 2018. More than half of US homes now subscribe to a streaming service, spending $2.1 billion a month. URL https://www.cnbc.com/2018/03/19/streaming-services-americans-spend-2-point-1-billion-a-month-in-55-percent-homes.html. Accessed: 2021-05-16.

Dai, H.; and Zhang, D. J. 2019. Prosocial Goal Pursuit in Crowdfunding: Evidence from Kickstarter. *Journal of Marketing Research* 56(3): 498–517.

Datta, H.; Knox, G.; and Bronnenberg, B. J. 2018. Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery. *Marketing Science* 37(1): 5–21.

Forbes. 2021. The Creator Economy, NFTs And Marketing. https://www.forbes.com/sites/kianbakhtiari/2021/04/18/the-creator-economy-nfts-and-marketing/. Accessed: 2021-05-10.

Gomez-Uribe, C. A.; and Hunt, N. 2016. The Netflix Recommender System: Algorithms, Business Value, and Innovation. *ACM Transactions on Management Information Systems* 6: 13:1–13:19.

Guardian. 2019. Is streaming killing the radio star? URL http://www.theguardian.com/media/2019/feb/08/is-streaming-killing-the-radio-star. Accessed: 2021-05-16.

Hair, L. 2021. Friends, not ATMs: parasocial relational work and the construction of intimacy by artists on Patreon. *Sociological Spectrum* 41(2): 1–19.

Hamilton, W.; Zhang, J.; Danescu-Niculescu-Mizil, C.; Jurafsky, D.; and Leskovec, J. 2017. Loyalty in Online Communities. *Proceedings of the International AAAI Conference on Web and Social Media* 11(1): 540–543.

Hogan, M. 2015. Upstream Effects of the Streaming Revolution: A Look into the Law and Economics of a Spotify-Dominated Music Industry Student Notes. *Colorado Technology Law Journal* 14(1): 131–152.

Holtz, D.; Carterette, B.; Chandar, P.; Nazari, Z.; Cramer, H.; and Aral, S. 2020. The Engagement-Diversity Connection: Evidence from a Field Experiment on Spotify. *arXiv:2003.08203 [cs]* .

James, T. L.; Shen, W.; Townsend, D. M.; Junkunc, M.; and Wallace, L. 2020. Love cannot buy you money: Resource exchange on reward-based crowdfunding platforms. *Information Systems Journal* .

Khan, M. L. 2017. Social media engagement: What motivates user participation and consumption on YouTube? *Computers in Human Behavior* 66: 236–247.

Mollick, E. 2014. The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing* 29(1): 1–16.

PWC. 2019. Getting personal: Putting the 'me' in entertainment and media. URL https://www.pwc.com/gx/en/entertainment-media/outlook-2019/entertainment-and-media-outlook-perspectives-2019-2023.pdf. Accessed: 2021-05-16.

Regner, T. 2020. Crowdfunding a monthly income: an analysis of the membership platform Patreon. *Journal of Cultural Economics* .

Waller, I.; and Anderson, A. 2019. Generalists and Specialists: Using Community Embeddings to Quantify Activity Diversity in Online Platforms. In *WWW '19*.

Wlömert, N.; and Papies, D. 2016. On-demand streaming services and music industry revenues — Insights from Spotify's market entry. *International Journal of Research in Marketing* 33(2): 314–327.