Signaling in Paid Product Placements: Theory and Evidence from Sponsorship Disclosure on Twitch

Ivan Li*

January 31, 2025

Abstract

Placing products in front of consumers frequently involves side payments. Such payments often require legal disclosure, but the style and the extent of such disclosure can be obfuscated. Nevertheless, prominent disclosure is sometimes chosen voluntarily. I collect novel data from Twitch.tv on revealed preferences for highly-visible (versus obfuscated) disclosure which let me study why prominent disclosure of sponsorships exists. I develop a theoretical signaling model that allows influencers to selectively use prominent disclosure to convey sponsor alignment, balancing immediate engagement revenue against potential future costs like credibility or reputational damage. I demonstrate that equilibria from this model can be empirically validated using a Wu-Hausman test by comparing the coefficients of OLS and IV regression. Empirical evidence from Twitch supports a partially-separating equilibrium where influencers visibly disclose well-aligned ("high type") sponsors but otherwise pool unaligned ("low type") sponsors with organic (non-sponsored) content to mitigate reputational damage from "sellout" effects. The signaling model can assess short-term responses to disclosure regulation and can be extended to product placement contexts outside of influencer marketing like radio, TV, and search platforms, explaining why equilibrium disclosure strategies differ across contexts.

Keywords: Sponsorship disclosure, Signaling, Product placement, Influencer marketing, Online livestreaming, Brand alignment

---

*Ivan Li is an Assistant Professor of Marketing at the Jindal School of Management, University of Texas at Dallas. Email: yilun.li@utdallas.edu. I am extremely grateful to my advisor Wesley Hartmann for terrific guidance. I thank Lanier Benkard, Ryan Dew, Sam Goldberg, Garrett Hagemann, Joe Hall, Yewon Kim, Samir Mamadehussene, Navdeep Sahni, Upender Subramanian, Ali Yurukoglu, and seminar participants at Carngie Mellon University, Stanford GSB, University of Delaware, UT Dallas, and Yale SOM for helpful comments. Researcher(s)' own analyses calculated (or derived) based in part on data from Twitch.tv's public API. The conclusions drawn from the Twitch API data are those of the researcher and do not reflect the views of Twitch or Amazon. All errors are my own. This paper comprises part of author's first dissertation chapter. Job Market Paper circulated as *Sponsorship Disclosure and Reputation in the Creator Economy: Evidence from Twitch.*

1 Introduction

Consumers frequently encounter situations where products are placed in front of them. In grocery stores, shelves are carefully organized to cater to consumer needs. Characters in the *Barbie* movie only drive Chevrolet cars. Social media platforms are flooded with product endorsements by celebrities and influencers. Some placements reflect the organic preferences of the store or celebrity, while others involve financial transactions that consumers may not be aware of. In certain contexts, disclosure of such paid placements is legally required but the party receiving payments can choose its degree of subtlety. Coase (1979) argues that disclosing such payments can be disruptive, as was the case with song placements on radio in the 1950's and '60s. However, if disclosure can be highly diminished to limit disruption in some contexts, why then would prominent disclosure ever occur?

Prominent disclosure of sponsored products can influence whether consumers perceive product placements as genuine or merely transactional. Such perceptions affect consumer purchasing decisions, trust, and reputations of those participating in the marketplace. Moreover, insights into disclosure practices can inform regulatory bodies aiming to protect consumers from deceptive advertising while balancing the interests of marketers and those accepting paid placements. Not all disclosure regulation is good; for example, payola regulation in the 1960s set back music discovery and the indie music scene as big record labels fought for the disclosure of pay-to-play songs on radio. Even though consumers enjoyed paid-for rock-and-roll music on the radio, the regulation killed their trust in popular radio DJs and removed a channel for smaller artists and labels to break into the industry (Coase, 1979). By analyzing which factors lead organizations or influencers to choose certain disclosure levels, we can gain insights into how disclosure norms emerge and how public policy might best protect consumers' interests while preserving the legitimate benefits of paid placements and sponsorships.

Disclosure behavior regarding paid placements has been challenging to study empirically because it is difficult to acquire data with appropriate variation. Variation in disclosure is commonly absent since disclosure is often regulated and required to be prominent. In some settings, non-disclosure is unobserved and requires complex algorithms and assumptions just to identify (Bairathi and Lambrecht, 2023; Ershov and Mitchell, 2023; Ershov et al., 2023). Furthermore, not many settings have repeated disclosure decisions. A panel of revealed preferences for disclosure by the payment-receiving party is necessary to separate preferences for disclosure from contexts favoring disclosure.

This paper studies why voluntary, prominent disclosure occurs in contexts when a party receiving payment for paid placements can choose the intensity of disclosure. I construct a theoretical signaling framework that generalizes across various product placement contexts, explaining why variation in disclosure practices exists. I characterize the conditions under which a nondisclosure pooling equilibrium occurs and the conditions under which a partially-separating equilibrium with voluntary disclosure occurs. Then, I provide proof that a common econometric test can be used to test for the empirical existence of the partially-separating equilibrium of my signaling model.

I apply this test to a unique dataset of content and disclosure choices comprising more than 1,000 English-speaking influencers ("streamers") on Twitch.tv, the world's largest online video game livestreaming platform. I focus on game developer (game dev) sponsorships, where developers pay streamers to play their video game live. The most unique aspect of my setting is that disclosure policies have been enacted but enforcement has been lenient, enabling streamers to “hide” disclosure labeling of sponsored streams behind a long string of text in their stream titles. I observe a panel of influencer stream decisions, which importantly includes within-influencer variation in game dev sponsored streams and in revealed preferences for disclosure, allowing me to identify influencers’ behaviors regarding disclosure.

The model is a two-player static game between an influencer (party receiving sponsorship/paid placement) and a follower (consumer). The influencer is exogenously assigned one of three types – high, low, or organic – detailing their “alignment” or match value with the sponsor. They can choose whether to prominently disclose the sponsorship (high disclosure) or to mitigate disclosure (low disclosure). High disclosure commands engagement that differs from low disclosure, but also incurs a “sellout effect” negatively impacting the influencer’s reputation or brand equity. Crucially, I assume these ”sellout effects” – which may materialize over an extended period of time – can be sufficiently captured by a reduced form cost function that varies with the influencer’s type and the disclosure level. This allows me to formulate the game with a structure similar to Spence (1973)’s signaling model.

The game proceeds with the following timing. The influencer observes their private type and makes a disclosure choice that maximizes their expected payoffs. The follower observes their own private cost of engaging and the influencer’s disclosure choice. They form beliefs about the influencer’s type, choose to engage or ignore the content, and then receive payoffs. If the follower engages, they reveal the influencer’s true type which affects their payoff. Next, the follower makes decisions about future interactions like unfollowing the influencer. The consequences of these actions with long-run effects are captured by the aforementioned reduced form cost function in the influencer’s payoff. Once the follower performs their actions, the influencer receives payoffs.

After characterizing the pooling and partially-separating equilibria of my model, I conduct theoretical counterfactual exercises to analyze the effects of a blanket high disclosure policy. Follower engagement with organic content improves if a partially-separating equilibrium had initially existed, but the effects on engagement are ambiguous and dependent on beliefs and relative magnitudes of the types if a pooling equilibrium had existed instead.

Two necessary conditions need to be empirically validated to support a partially-separating equilibrium. One key condition is a *single-crossing condition*, which states that influencers’ indifference curves can only cross once. This condition is satisfied if the cost of disclosure is increasing in the disclosure level and is increasing more for low types than high types. I demonstrate that under some assumptions, a Wu-Hausman test for endogeneity – comparing regression coefficients from ordinary least squares (OLS) and instrumental variables (IV) regressions – is a valid test for single-crossing. The second condition requires that the probability of engaging under high disclosure must be higher than the probability of engaging under low disclosure. Regression coefficients from the OLS specification enable me to assess this condition.

I evaluate both conditions using the Twitch data. Conditional on a stream being sponsored by a game developer, high disclosure occurs only 14% of the time. Yet, OLS regression results show that high disclosure streams are correlated with 5-12% higher viewership compared to low disclosure streams, commanding viewership at or above viewership of organic streams. These two facts are consistent with characteristics of a partially-separating equilibrium. The IV regression returns a significant and negative impact of high disclosure on reputation, and the Wu-Hausman test shows that the IV regression coefficient of interest is significantly different from the OLS coefficient, satisfying the single-crossing condition. The instrument I use proxies a unique feature of the industry; certain sponsors may include clauses in their contracts with streamers that mandate high disclosure for various reasons such as fear of regulation enforcement or other idiosyncratic preferences. This instrument effectively allows me to observe off-path equilibrium outcomes as some low types are forced to disclose by a sponsor's directive. Moreover, the IV regression finds a positive effect of high disclosure streams on viewership which is consistent with off-equilibrium path outcomes from the signaling model.

I construct a measure of alignment to provide more evidence that high types are the ones who select into high disclosure. This is a correlational measure between the qualitative characteristics of video games and a streamer's historical frequencies of playing games with these characteristics. For a subset of streamers whose historical preferences are relevant to their game choice, streamers choose to disclose games that are more aligned. The qualitative characteristics of high disclosure sponsored games looks very similar to those of organically chosen games. The games that are not prominently disclosed are much lower in alignment and look much different than a typical organically chosen game. This finding further supports the partially-separating equilibrium outcome where disclosure is only chosen by high types.

Lastly, I use a staggered event study to assess the context surrounding sponsorship and disclosure events (Borusyak et al., 2024). Pre-trends can assess the selection into disclosure, and post-trends can assess the dynamic effects of reputation costs. I find that influencers who choose high disclosure have significantly higher organic content viewership in the weeks leading up to the sponsored event than those who choose low disclosure. Furthermore, their viewership remains higher post-sponsorship. If influencers prominently disclosed low-types, or if disclosure was assigned randomly, then significant differences in pre and post-trends should not exist.

Related Literature. My main contribution is showing that influencers use disclosure as a mechanism to signal their alignment with a paid product placement to followers. My theoretical signaling framework can be extended to understand disclosure or non-disclosure behavior other contexts with product placements, like radio (Coase, 1979) and search platforms (Sahni and Nair, 2020a). Within influencer marketing, the model generalizes sponsorship disclosure results from Karagür et al. (2022), Ershov and Mitchell (2023), and Bairathi and Lambrecht (2023) on Instagram. This builds on the discussion of product placements in more traditional marketing settings like slotting fees (Sullivan, 1997; Sudhir and Rao, 2006; Hristakeva, 2022), radio (Coase, 1979), and television (Russell, 2002).

I further contribute to the burgeoning literature on influencer marketing. Much of this literature is purely theoretical, especially regarding disclosure (Berman and Zheng, 2020; Fainmesser and Galeotti, 2021; Mitchell, 2021; Pei and Mayzlin, 2022). This literature tends to focus on mechanisms of consumer demand (e.g. word-of-mouth) and consumer welfare. On the supply side, Nistor et al. (2024) provides a theoretical framework to understand when influencers switch from producing organic to sponsored content. I abstract away from influencers' growth and solely focus on influencers' disclosure decisions.

I also measure viewership of sponsored content, incorporating brand-influencer alignment into the viewership demand equation. This provides advertising brands another factor to consider when assessing the potential effectiveness of influencer marketing campaigns (Rajaram and Manchanda, 2020; Morozov and Huang, 2021; Li et al., 2021; Yang et al., 2021; Hofstetter et al., 2023; Nistor and Selove, 2023). I define and quantify "brand alignment" between a brand and an influencer using historical revealed preferences, providing future researchers a method to quantitatively analyze hypotheses regarding influencer credibility, authenticity, and influencer-product congruence (Avery and Israeli, 2020; Schouten et al., 2020; Kim and Kim, 2021; Li et al., 2021; Pöyry et al., 2021; Amano et al., 2023). While Cheng and Zhang (2022) consider reputation burning and brand-influencer fit, I show that disclosure is a crucial lever used by influencers to signal good fit and to mitigate reputational costs. My long panel of influencer choices allows me to identify behaviors regarding disclosure that lab experiments are unable to consider (Boerman, 2020; Kay et al., 2020).

Moreover, I contribute to the literature on native advertising and its disclosure (Evans et al., 2019; Aribarg and Schwartz, 2020; Sahni and Nair, 2020b). One paper to note is Sahni and Nair (2020a), who use a field experiment to study effects of disclosing search ads on a Yelp-like restaurant platform. They find that disclosure increases clickthrough and calls to advertising restaurants. They attribute this to a "signaling effect," whereby customers perceive advertising restaurants to be higher in quality than non-advertising ones. My findings are consistent in that disclosure seems to serve as a signal in some circumstances. However, I find that disclosure can be detrimental especially when reputational costs of disclosure are high. These circumstances occur when the sponsor is not well aligned with the influencer.

Finally, my paper is one of the few that is set in the online livestreaming industry. There is some focus on non-sponsorship mechanisms of influencer monetization such as donations and various pay-what-you-want mechanisms (Lin et al., 2021; Lu et al., 2021). Morozov and Huang (2021) study the effects of streaming on video game usage more generally, treating all video game streaming as advertising. Simonov et al. (2021) uses a specific subset of Twitch data from streams of Counter-Strike:Go tournaments in addition to viewer-level chat data to study the role of suspense. I add to this literature by focusing on influencers' revealed preferences for disclosure and the role of disclosure as a high alignment sponsored content signal.

2 Institutional Detail and Data

2.1 The Online Livestreaming Industry

The online livestreaming economy has been booming in recent years. Audiences watched almost 8.5 billion hours of online livestreams in Q2 2024[1]. The most popular livestreamers command tens if not hundreds of thousands of concurrent viewers and sign exclusive streaming contracts worth tens of millions of dollars [2]. Twitch, specifically, occupies about 60% market share. On average, there are 2.5 million concurrent viewers on Twitch.tv and 90,000 unique live streamers at any moment.[3] Influencers on the Twitch platform usually stream themselves playing video games or "just chatting," which is a general category for non-gaming related or "in real life (IRL)" streams.

Streamers have three broad ways to monetize. The first way involves Twitch-embedded ads, which are pre-negotiated by Twitch and its advertisers. Similar to Youtube ads, these ads usually run when a consumer initially lands on a streamer's livestream. Streamers can also press an "ad button" whenever they want to run such ads. There is no way to obfuscate these ads, payment depends upon the calculated reach of the ads, and in recent years, these ads have become unblockable and unskippable. The second way involves direct contributions from viewers. Viewers can unlock a streamer's premium channel features by becoming a paid "subscriber," which costs anywhere from $5 to $25 a month. Streamers then receive a portion of the subscription revenue. Streamers can also receive donations from viewers through Twitch or a third party.[4]

External sponsorships are the third way, but even then there are nuances. There exists two subcategories of sponsorships - brand deals and game developer deals. Game developer (game dev) deals are product demonstrations or game playthroughs that alter alter the content of a stream. The typical game developer deal involves a streamer playing a sponsored game for a few hours. A brand deal generally involves any product that is not a video game itself. Apparel, computer hardware, and food delivery services are examples of brand deals. This paper only considers game dev deals; brand sponsorships are ignored because they generally do not fundamentally alter the content of the stream.

Sponsorships require negotiation around compensation and deliverables. Disclosure, according to industry insiders, is rarely part of the negotiation. Some sponsors may have certain preferences over disclosure practices,[5] but sponsors usually do not dictate how influencers should disclose.

2.2 Data

My data comes from two main sources. Streaming data is collected from Twitch.tv's API. I collect data on the top 430 english speaking Twitch streamers starting in February 2021. In August

[1] Stream Hatchet Live Game Streaming Trends Q2 2024

[2] Anecdotal evidence from streamers within the industry, see: `https://www.youtube.com/watch?v=qDMJQeHxYeQ`

[3] See: `https://www.twitchtracker.com`

[4] Livestream donations are the object of focus in Lin et al. (2021) and Lu et al. (2021)

[5] For example, in my data, the game *Legends of Runeterra* almost always has #ad at the beginning of the stream title across different influencers

2021, I expanded the data collection to the top 1,300 english speaking Twitch streamers. In this paper, the data collection period ends on April 30, 2023.

Every 5 minutes, I am able to obtain, for each streamer, the live/offline status of their stream, the number of concurrent viewers (if live), the number of total views a channel has, the title of the stream, the game being played, and the number of users following the channel.

Certain metrics are not updated every five minutes, so I aggregate data up to the user-stream-game level. I drop stream-game combinations that are live for less than 30 minutes. For example, if a user is live for 6 hours on Sunday, October 17th and they spend their first 2 hours streaming League of Legends, the next 1 hour 45 minutes streaming Grand Theft Auto V, the next 15 minutes "Just Chatting", and finally spend their last 2 hours going back to League of Legends, this one stream session would be broken up into three observations in my data even though there are four stream-game combinations.

Statistic (per streamer)	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Observations	579.58	432.96	10	329	523	746	7,890
Num. streams	384.69	197.44	9	250	391	512	1,841
Num. unique games	51.73	77.79	1	12	27	61	1,148
Num obs any sponsor	25.53	59.06	0	1	9	28	1,110
Num obs any hi disc sponsor	2.14	6.49	0	0	0	2	95
Num obs game dev	10.82	21.88	0	0	4	13	356
Num obs game dev hi disc	1.49	4.96	0	0	0	1	74
Avg. conc. viewership	3,683.59	8,476.63	94.80	739.59	1,458.85	3,037.49	105,018.90
Avg. stream+game length (hr)	4.83	2.26	1.28	3.22	4.52	5.91	23.38
Game dev sponsor	0.017	0.031	0.000	0.000	0.008	0.022	0.383
Any sponsor	0.040	0.067	0.000	0.003	0.019	0.049	0.848
Game dev hi disc pct (conditional on ad)	0.145	0.243	0.000	0.000	0.000	0.222	1.000
Any sponsor hi disc pct (conditional on ad)	0.111	0.207	0.000	0.000	0.000	0.125	1.000
Current followers	621,062	1,133,851	5,934	133,334	286,355	638,654	17,807,250
Initial followers	461,453	942,332	1,571	99,808	200,206	456,120	16,714,288
Follower change	159,609	331,725	−53,903	14,603	45,111	152,898	3,832,885

Table 1: Streamer summary statistics, 1159 streamers

After selecting streamers based on some criteria (see Section A.1 for more detail), I am left with 1,157 streamers and around around 670,000 observations at the user-stream-game level. Streamer-level summary statistics are provided in Table 1. Over two-thirds (821) of all streamers have done a game dev sponsored stream. Out of these 821 influencers, 377 have highly disclosed a sponsored stream at least once. The median streamer has 523 stream-game observations over 391 streams, about 1.4 separate game sessions per stream. The median streamer has 4 observations with a game dev sponsor, comprising about 1% of their observations.

The median streamer commands just under 1,500 average concurrent viewers (ACV) during a stream.[6] ACV is defined as the mean number of unique viewers on a stream at any point the streamer is live. This is a crucial metric, as ACV is how Twitch culture tends to measure the size of a streamer.[7]

[6] These are big streamers; for example, `https://twitchtracker.com/day9tv` is a $\sim 1,500$ viewer streamer who is in the top 0.03% of Twitch

[7] See `https://www.quora.com/Why-do-Twitch-streamers-refer-to-each-other-as-Andy` as an example

Game dev sponsored content streams occur rather infrequently; about 12,000 or 1.7% of all observations are sponsored. About 1,600 of these are considered "prominently disclosed" under my definition (see Section 2.3). At the beginning of my data sample period, sponsored livestreams on Twitch occupied 3% of total watch time.[8]

Game characteristics are collected from the Internet Game Database (IGDB) API, which is a website owned and operated by Twitch. For each game, I can access characteristics such as its genres, themes, storylines, release date, user and critic ratings, and much more. Twitch uses IGDB on its own website to make it easier for viewers to search for games. In my data sample, streamers play almost 9,000 unique video games.

2.3 Identifying Sponsored Streams

One implicit assumption I make is that streamers truthfully disclose all sponsored content. There is good reason to believe that disclosure happens; FTC regulations require disclosure of any "material connections" between an influencer and a brand[9], and so does Twitch's terms of service[10]. Streamers in my data are among the most popular on Twitch, many of whom treat streaming as a full-time job. The threat of enforcement from the FTC and Twitch to their livelihoods should be enough to ensure disclosure [11]. The ability to obfuscate while complying should also limit non-disclosure. Discussions with talent management agencies in this industry support this claim that streamers generally are well-behaved with respect to disclosing sponsored content.

When viewers browse for a stream, they can see a thumbnail picture of the livestream, as well as information such as the title of the stream, the name of the streamer, and the game being played currently by the streamer. Figure 1 shows what a follower observes when browsing for a stream. Prior to clicking on a stream channel and watching the stream, any potential viewer can only find out about the sponsored nature of the stream through the stream title.

I identify sponsored content using a simple string match on the stream titles. Within the stream titles, I search for instances of #ad, #sponsored, and variations of #*partner (e.g. #EpicPartner). Every stream that simply contains one of these hashtags is tagged as potentially sponsored. To distinguish game dev deals from brand deals, I manually look for each observation if the name of the game being played is contained in the stream title of the sponsored content. Some examples of titles related to brand deals and game dev deals are provided in Table 2. I separately identify high disclosure and low disclosure using the location of the hashtag in the stream title. The length of the stream title dynamically adjusts depending on the screen resolution of the viewer's device. The typical length displayed on the screen is between 20 and 40 characters. I define high disclosure as an indicator function taking on the value 1 if the start of the hashtag is located within 15 characters

[8] Stream Hatchet Q1 2021 report

[9] https://www.ftc.gov/tips-advice/business-center/guidance/ftcs-endorsement-guides-what-people-are-asking

[10] https://www.twitch.tv/p/en/legal/terms-of-service/

[11] Teami Detox Teas is an example of a company recently fined in 2020; celebrities endorsing the product such as Cardi B were warned for their lack of disclosure: https://www.ftc.gov/news-events/press-releases/2020/03/tea-marketer-misled-consumers-didnt-adequately-disclose-payments

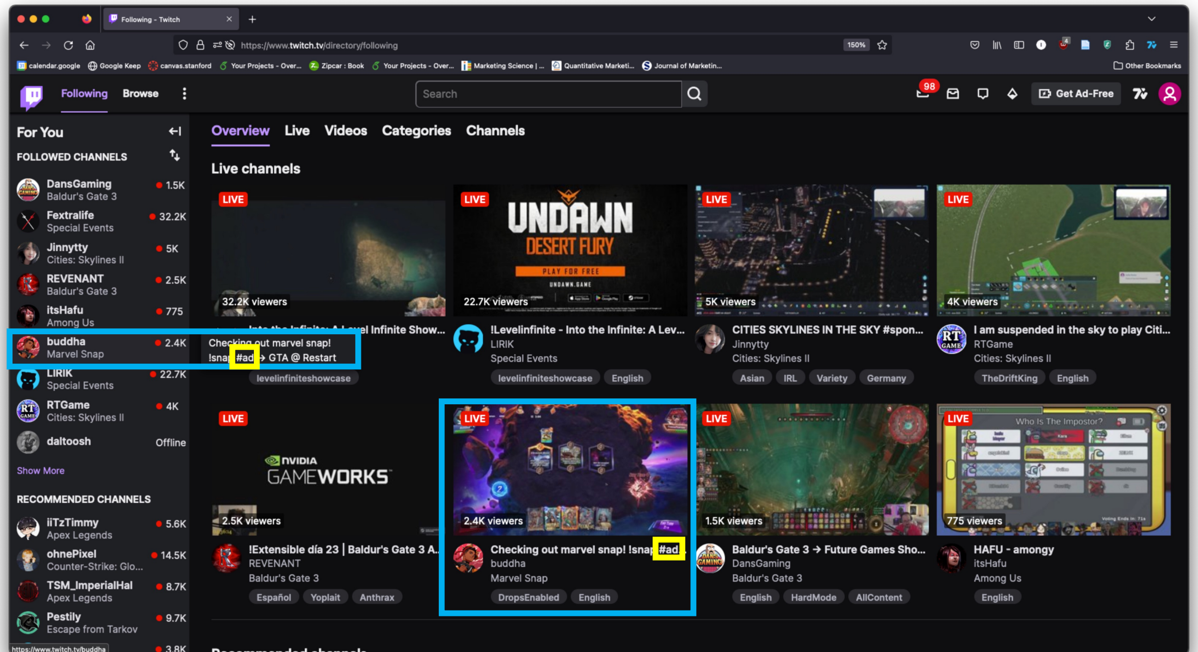

Figure 1: Example of Twitch following homepage for a follower of the streamer Buddha. The streamer (outlined in blue) can appear on the sidebar and on the main panel. Disclosure (outlined in yellow) is only visible in the stream title in both places.

	High Disclosure	Low Disclosure
Game dev	#Sponsored: Legends of Runeterra (Game indicated on stream: Legends of Runeterra)	ROCKET LEAGUE THEN Marvel Strike Force #ad !marvel (Game indicated on stream: MARVEL Strike Force)
Brand	#Sponsored by Universal — Follow @shroud on socials (Game indicated on stream: Apex Legends)	Herman Miller Gaming Giveaway !hmgaming #sponsored (Game indicated on stream: Battlefield 2042)

Table 2: Examples of stream titles

from the front of the stream title. My empirical results in this paper are robust to alternative definitions of high disclosure, including arbitrary locations greater or less than fifteen characters from the beginning of the stream title (See appendix A.6). Table 2 gives examples of high and low disclosure ads.

The location of the disclosure label for game dev deals is displayed in Figure 2. The red dashed line indicates the 15th character, where I set my cutoff for high disclosure. There is a mass at zero, indicating that a large number of sponsored streams have the hashtag immediately at the beginning of the stream title. There are no other large masses that jump out, indicating that the decision to put the hashtag at the beginning might be selective.

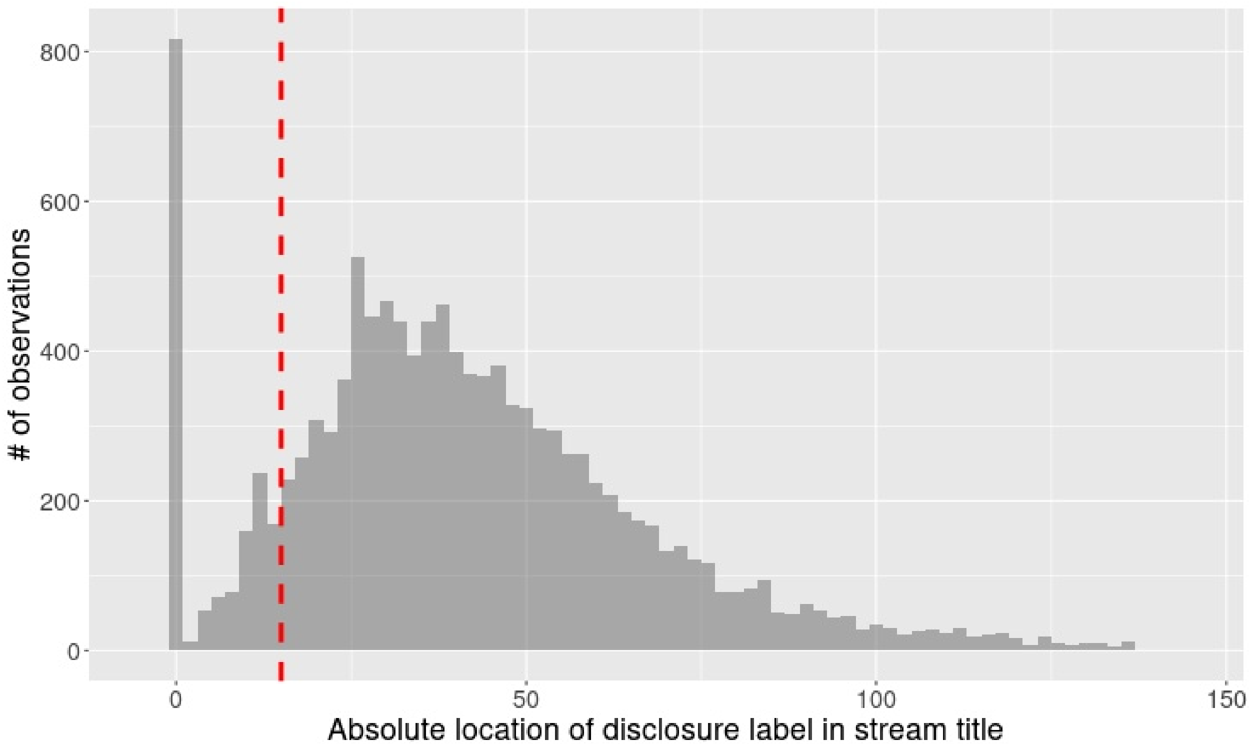

Figure 2: Absolute location of advertising disclosure, game dev deals only

3 A Stylized Model

3.1 Model Setup

The game is set up as follows. There are two players: one "influencer" i and one "follower" f. In general, the influencer can be thought of as any party receiving payment for paid placement, and the "follower" is any consumer who engages downstream with this party.[12] Before any actions occur, i exogenously realizes a paid product placement (sponsor) opportunity with alignment θ, which is unobserved to follower f. There are three discrete alignment types $\theta \in \{\theta_H, \theta_L, \theta_0\} \in \mathbb{R}$,

[12]For example, i can be a TV program and f can be a viewer.

corresponding to high, low, and organic alignment respectively.[13] Followers treat alignment as a vertical characteristic given an influencer-brand pair. Conceptually, θ is a match value which can vary horizontally across influencers for a fixed brand/sponsor. That is, different influencers may have a different alignment type for the same sponsor.[14]

Conditional on realizing one of the sponsored types θ_H or θ_L, these types have conditional probabilities $p_\theta \in \{p_H, p_L\}$ which individually exist on $(0, 1)$ and collectively sum to 1. Furthermore, let $q \in (0, 1)$ be the probability that a sponsorship exists, such that $p_\theta \cdot q$ is the true incidence of sponsored content of alignment θ and $1 - q$ is the incidence of θ_0. Both p_θ and q are common knowledge. I place the restriction on the types $\theta_H > \theta_L$, forcing the high type to be strictly better than the low type, and $\theta_0 > \theta_L$. I take no stance on comparing θ_H and θ_0.

In the first stage, the influencer has two actions to choose from, $j \in \{LD, HD\}$, which map to low and high disclosure respectively. The influencer chooses j that maximizes their expected payoff, $\mathbb{E}[\pi_{ij}]$. Depending on context, LD can be diminished disclosure or the absence of disclosure; the θ_0 type is restricted to only choosing LD. HD can be prominent disclosure or the presence of disclosure. Assume that the influencer's outside option of no content does not exist in this setting and that the influencer cannot switch between sponsored content and organic content freely; for example, a contract has been signed and the product placement must occur. This also allows me to focus on just the disclosure decision and abstract away from the choice to create sponsored content and payments to the influencer, which is private contract data and often unobserved.

Once the disclosure choice j has been made, follower f observes j and makes a choice $k \in \{engage, ignore\}$ that maximizes their expected utility $\mathbb{E}[u_{fk}]$. This expectation is affected by their posterior belief of θ after observing j, which I denote $\mu_{\theta j} = Pr(\theta|j)$. Beliefs have sums $\sum_{\theta \in \{\theta_H, \theta_L\}} \mu_{\theta,HD} = 1$ and $\sum_{\theta \in \{\theta_H, \theta_L\}} \mu_{\theta,LD} < 1$. The latter inequality reflects the possibility that LD sponsored content can be confused for θ_0 non-sponsored content by consumers.

If $k = engage$ is chosen, the follower incurs a cost g drawn from a distribution $\mathcal{G}(\cdot)$ with support on the non-negative real numbers. This cost g is private information to the follower and can be interpreted as an opportunity cost to keep looking for content. Engaging fully reveals the true alignment of the content, θ, and the follower receives $\theta - g$ as their payoff. If $k = ignore$ is chosen, the follower receives zero payoffs. In summary, the follower's payoffs are:

$$u_{fk} = \begin{cases} \theta - g & \text{if } k = \text{engage} \\ 0 & \text{if } k = \text{ignore}. \end{cases} \tag{1}$$

If the follower engages with the content, they also choose the "intensity" of engagement, $v(\theta)$, where v is an increasing function in θ. This is a static, reduced form way of modeling viewing time, liking the content, etc., which enters i's payoff. The follower simultaneously makes decisions

[13] Alignment does not necessarily entail vertical quality. For example, any sports drink may be well-aligned with an athlete but any cryptocurrency sponsor is likely poorly-aligned

[14] For the sake of exposition, I will refer to the realization of the brand-influencer alignment as an influencer's "type"

that affect their future engagement with the influencer. In practice, a follower can unfollow an influencer if they do not like sponsored content. The implication of unfollowing in a dynamic model with multiple periods would be a decrease in f's engagement (and therefore i's payoffs) in future periods. Followers' beliefs can also update and impact future decisions. For tractability, I model these decisions with dynamic effects as a reduced form reputation/brand equity cost $c(j|\theta) : \{HD, LD\} \to \mathbb{R}_{\geq 0}$, where the cost incurred by influencer's action j depends on the alignment θ. These costs are scaled by the engagement intensity v. A cost of disclosure is incurred even if the follower doesn't engage with the content because the sponsored nature of the content is visible due to the influencer choosing HD.

After the follower makes their decisions, the influencer receives payoffs for their choice j:

$$\pi_{ij} = \begin{cases} 1\{k = engage\}v(\theta)(1 - c(HD|\theta)) - 1\{k = ignore\}c_{HD} & j = HD \\ 1\{k = engage\}v(\theta)(1 - c(LD|\theta)) & j = LD \end{cases} \tag{2}$$

The simplification of the dynamic effects of influencer's disclosure choice j into a reduced form cost function $c(j|\theta)$ turns the disclosure game into a signaling game akin to Spence (1973). I impose two assumptions on the influencer's cost of disclosure, $c(j|\theta)$. First, a monotonicity assumption implies $c(HD|\theta) > c(LD|\theta)$ for all θ. Practically, drawing attention to the sponsored nature of the content incurs a "sellout" effect, more negatively impacting reputation.[15] Second, $c(LD|\theta_H) = c(LD|\theta_L) = c$. This restriction states that the costs of low disclosure are independent of sponsor alignment.[16]

The cost $c_{HD} \leq \min_\theta v(\theta)c(HD|\theta)$ is the cost of high disclosure without engagement. The inequality is justified because on many platforms, followers can easily unfollow on the page where they engage with the content. If they ignore, followers have to go out of their way like clicking into the influencer's profile page in order to unfollow. This takes extra effort which may dissuade followers from unfollowing, so high disclosure is less costly when ignored.

The extensive form of the game is displayed in Figure 3. To reiterate the timing of the model:

1. i observes θ and chooses j, taking expectations over f's decision.
2. f observes j and g, and chooses k, taking expectations over θ.
   (a) if $k = engage$, f engages with intensity $v(\theta)$ and receives payoffs
   (b) if $k = ignore$, f receives zero payoffs.
3. f makes decisions about future interactions (including updating beliefs) with i: $c(j|\theta)$
4. i receives payoffs.

The equilibrium concept I will use for this model is perfect Bayesian equilibrium (PBE). This equilibrium is characterized by: i.) an action j for each type $\theta \in \{\theta_H, \theta_L, \theta_0\}$, ii.) Followers' prior

[15]Some may argue that "honesty" or "transparency" diminish sellout or negative reputational effects. Evidence about this is mixed and studies leveraging exogenous variation usually find negative effects of disclosure (Karagür et al., 2022; Ershov and Mitchell, 2023), indicating that this assumption is plausible.

[16]I discuss this assumption in Appendix A.5.

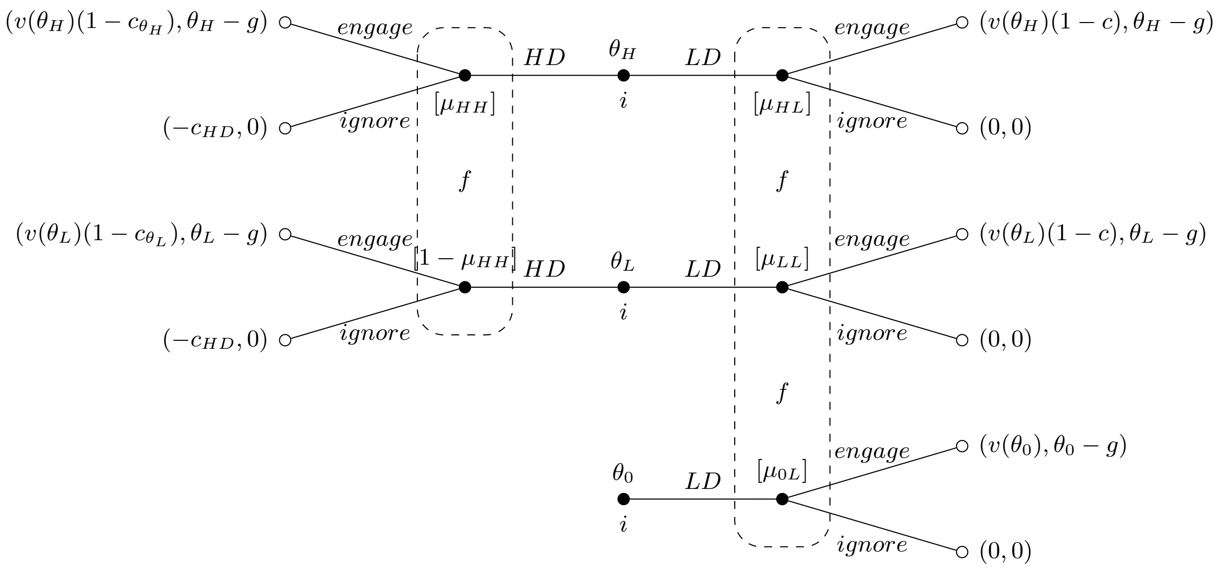

Figure 3: Extensive form of the signaling game. To simplify notation, I denote $c(HD|\theta_H) = c_{\theta_H}$ and $c(HD|\theta_L) = c_{\theta_L}$. The notation for beliefs is also shortened to fit the graphic. For example, $\mu_{HH} = \mu_{\theta_H,HD}$.

beliefs over θ, iii.) Followers' posterior beliefs over θ after observing j, $\mu_{\theta j}$. The equilibrium must satisfy optimality of choices for both i and f, and all beliefs must be consistent with Bayes' rule.

3.2 Pooling Equilibria

In many marketing settings, we observe no disclosure of paid placements whatsoever. I study the conditions of this pooling equilibrium first.

Lemma 1 *At least one pooling equilibrium where $j = LD$ for every θ exists under the assumptions made in Section 3.1.*

See Appendix A.3 for the proof. We observe the $j = LD$ pooling equilibria for product placements in settings such as TV shows, movies, and radio. High disclosure costs $c(HD|\theta)$ are large and in many circumstances $\mu_{\theta_H,HD}$ is small, making it difficult for $P(engage|HD)$ to be larger than $P(engage|LD)$. As a result, there is little incentive for types (especially θ_H) to deviate from low disclosure.

For TV shows and movies, the costs of disclosure are large because disclosing might be visually distracting, degrading the perceived quality of the show overall. In Coase (1979)'s radio example, the costs for the radio DJ were extremely large for two reasons. First, there was massive reputational damage from being accused as a "sellout." Radio shows of payola-affected DJs were cancelled or became much less popular after congressional trials exposed the existence of pay-for-play arrangements – even when there wasn't enough proof to implicate DJs for accepting undisclosed payments. Second, disclosure forces narrative breaks in the music which sound like advertisements, making the radio show much less palatable to listeners. Even though the sponsored music was of high quality

– genres like rock-and-roll were popularized by payola – the public backlash from payola hearings made it obvious why DJs did not want to disclose.

3.3 Partially-Separating Equilibria

On Twitch, I observe influencers voluntarily choosing both high and low disclosure (Table 1) which cannot exist in a low disclosure pooling equilibrium. One equilibrium outcome that contains both choices is a partially-separating equilibrium. In this equilibrium, at least one type remains distinguishable while the other two types pool on the same action. This setup allows the follower to perfectly identify one type and partially update beliefs when seeing the action chosen by the pooled types.

Lemma 2 *At least one partially-separating equilibrium where $j = HD$ for θ_H and $j = LD$ for all other θ exists under the assumptions made in Section 3.1*

The proof is in Appendix A.4, and I will explain the intuition below:

When θ_H types choose HD, they signal their θ_H alignment to f, increasing the probability of engagement. θ_L types are better off choosing LD because their cost of high disclosure, $c(HD|\theta_L)$, is too large. The boost they get from an increased probability of engagement does not offset their cost. A key condition for the existence of this equilibrium is the *single-crossing condition*, which in this model is $c(HD|\theta_H) < c(HD|\theta_L)$.

The "opposite" partially-separating equilibrium where θ_L chooses $j = HD$ and θ_H pools with θ_0 cannot exist because $P(engage|HD)$ is guaranteed to be lower than $P(engage|LD)$ due to assumptions on θ and the monotonicity of $G(\cdot)$. This violates one of the necessary conditions for a partially-separating equilibrium. A third partially-separating equilibrium where θ_H and θ_L pool on HD is discusssed in Appendix A.4.1. In short, this equilibrium is unlikely to occur in practice because of required conditions on θ.

My context of online livestreaming seems to exist in the environment of a partially-separating equilibrium. Reputation can be proxied in this setting with follower count, so $c(j|\theta)$ can be thought of as a change in followers.[17] In practice, there are two mutually exclusive types of "followers" who might engage with the influencer – actual followers and outsiders (non-followers). When a sponsor is type θ_H, streamers sometimes disclose to signal to outsiders that they are popular or reputable, attracting more engagement from this group and converting some into followers. Current followers can interpret high disclosure that the sponsor is a good match for them, driving up their probability to engage.

In contrast, if a sponsor type is θ_L and streamers choose HD, streamers reveal that they are willing to accept any sponsor for cash. Both followers' and non-followers' perception of the streamers may worsen - this is the "sellout" effect.[18] Non-followers who engage with θ_L content do

[17]See Cheng and Zhang (2022) for a similar interpretation of subscribers as reputation on Youtube.

[18]Some influencers may always disclose in order to be "honest" or "transparent." In our empirical analysis, influencer fixed effects will control for this behavior.

not convert to becoming followers at a high rate. Followers who engage may believe that streamers will continue to sell their stream time to the highest bidder and unfollow. Negative responses to the low type nature of the content could still occur if the streamer chooses LD, but fewer followers will realize that content is sponsored, mitigating "sellout" effects. Hence, the cost of disclosure for streamers is higher for θ_L types because sellout effects exacerbate negative beliefs about future content and is lower for θ_H types because positive signaling effects mitigate the sellout effects.

3.4 Counterfactual analysis: disclosure regulation

The topic of disclosure naturally leads to the question of regulation. What happens when regulation forces high disclosure? My model can provide limited analytical insights regarding counterfactual behavior under pooling and partially-separating equilibria.

3.4.1 Regulation of a low disclosure pooling equilibrium

If regulation occurs in a setting where a low disclosure pooling equilibrium exists, the counterfactual outcomes depend on the prior beliefs of the follower, μ, and the relative values of θ. We expect $\mu_{\theta_H,HD} = p_H$ and $\mu_{\theta_L,HD} = 1 - p_H = p_L$ as beliefs update post-regulation. For probability of engagement to decrease, we need:

$$\begin{aligned}
q(p_H\theta_H + p_L\theta_L) + (1-q)\theta_0 &\geq p_H\theta_H + p_L\theta_L \\
q(p_H\theta_H + (1-p_H)\theta_L) + (1-q)\theta_0 &\geq p_H\theta_H + (1-p_H)\theta_L \\
qp_H(\theta_H - \theta_L) - p_H\theta_H + p_H\theta_L &\geq \theta_L - q\theta_L - (1-q)\theta_0 \\
(q-1)p_H(\theta_H - \theta_L) &\geq \theta_L - q\theta_L - (1-q)\theta_0 \\
(q-1)p_H(\theta_H - \theta_L) &\geq (1-q)(\theta_L - \theta_0) \\
p_H &\leq \frac{\theta_0 - \theta_L}{\theta_H - \theta_L}, q \neq 1
\end{aligned}$$

If $\theta_0 > \theta_H$, then engagement must decrease post-regulation. If $\theta_H > \theta_0$, then engagement may increase or decrease depending on the proportion of high alignment sponsors in the population.

The above analysis is valid for short-run counterfactual outcomes. The model does not include an outside option of no stream and does not allow switching between organic and sponsored content, limiting our ability to study the long-run supply of sponsored content. To illustrate, p_H and θ_H were likely large in a setting like Coase (1979). The model would predict that payola regulation should have increased engagement for pay-to-play songs. In reality, $c(HD|\theta_H)$ was so large that DJs stopped accepting pay-to-play radio songs entirely. Consequently, over the long run the supply of sponsored content decreased. Even in a case like this where the reason for low disclosure pooling is likely known, the model does not reach the correct conclusions about long-run supply of sponsored content post-regulation without additional assumptions.

3.4.2 Regulation of a partially-separating equilibrium

When regulation occurs to a partially-separating equilibrium, some counterfactual outcomes are clear. Engagement for organic content increases for certain as θ_L types are now forced to pool with θ_H types. This should be true in both the short and long term. Sponsored content engagement levels should decrease in the short term because θ_L content now pools with θ_H content. Whether or not followers benefit from this regulation in the short-run depends on the relative values of θ_H and θ_0.

If $c(HD|\theta_L)$ is large enough, then in the long-run we should expect eventually lower supply of θ_L type content because the costs of disclosure hamper low alignment sponsorships payoffs. As mentioned previously, the lack of an outside option makes this intuitive argument difficult to formalize.

4 Empirical Results

In this section, I present descriptive evidence to test predictions made by the theory model, providing support for a partially-separating equilibrium (Section 3.3) where influencers disclose to signal highly aligned sponsors. I show that sponsored game dev streams are correlated with lower average concurrent viewership (ACV) and more negative follower change compared to organic content. Lower ACV is a short term cost that decreases the Twitch-ad and donation revenue for a particular stream, and fewer followers decreases future viewership, leading to less Twitch and sponsor incomes in the future. However, if a high-type sponsor chooses high disclosure, signaling effects improve ACV. I use an instrumental variables (IV) regression to recover unbiased estimates of the effects of disclosure. I then demonstrate how to use a Hausman test to empirically test the existence of a partially-separating equilibria. Next, I construct a measure of "brand alignment" using qualitative video game data and streamers' histories to show that streamers select into high disclosure when video games are well-aligned. Finally, I provide additional evidence for selection into high disclosure by running an event study that shows that streamers who choose high disclosure have higher ACV in the weeks leading up to the sponsored stream and that this effect persists post-sponsorship.

4.1 Viewership trends are consistent with a partially-separating equilibrium

I run a OLS regression with ACV and follower change as my dependent variables of interest on stream-game level data. Influencer i at stream observation t realizes the outcome Y_{it} according to the specification:

$$Y_{it} = \beta_0 + \beta_a ad_{it} + \beta_d HD_{it} * ad_{it} + \beta_x x_{it} + \nu_i + \tau_t + \xi_{dev} + \varepsilon_{it} \tag{3}$$

where ad is the indicator for a sponsored game dev stream; HD_{it} is high disclosure of the stream; x_{it} includes time-varying observable characteristics of streams, games, and influencers; ν_i is an influencer fixed effect; τ_t are month-year and day-of-week related fixed effects; and ξ_{dev} are game developer fixed effects. OLS regression measures equilibrium outcomes because it takes equilibrium actions, like HD_{it}, as given.

	Full sample		Game dev sponsors only	
	Log ACV	IHS Follower Change	Log ACV	IHS Follower Change
Game dev sponsor	−0.058	−0.597		
	(0.016)	(0.091)		
Game dev spon hi. disc.	0.114	0.144	0.050	−0.026
	(0.036)	(0.172)	(0.017)	(0.089)
IHS game age	−0.022	−0.031	−0.011	−0.030
	(0.002)	(0.007)	(0.003)	(0.016)
Same-week streams	−0.003	0.074	0.022	0.204
	(0.003)	(0.012)	(0.005)	(0.028)
Log Stream Length	0.199	0.947	0.081	0.817
	(0.005)	(0.023)	(0.015)	(0.070)
Drops	0.223	0.832	0.114	1.100
	(0.019)	(0.069)	(0.032)	(0.147)
Tournament	−0.029	0.035	−0.018	0.891
	(0.016)	(0.042)	(0.040)	(0.244)
Championship	0.143	0.231	0.149	0.028
	(0.052)	(0.100)	(0.209)	(0.678)
Giveaway	0.053	0.293	0.073	0.584
	(0.014)	(0.054)	(0.026)	(0.135)
Charity	0.029	0.010	−0.001	−0.418
	(0.021)	(0.078)	(0.054)	(0.335)
Subathon	0.100	0.032	0.293	0.605
	(0.034)	(0.082)	(0.069)	(0.268)
First Game	−0.223	0.096	−0.279	−0.276
	(0.009)	(0.027)	(0.020)	(0.220)
Log total followers	0.704	0.600	0.585	0.026
	(0.066)	(0.154)	(0.128)	(0.672)
Alignment	0.447	1.737	0.035	−0.031
	(0.035)	(0.140)	(0.058)	(0.332)
Num. obs.	668131	668131	11989	11989
R^2	0.872	0.586	0.907	0.605
Game Characteristics	Y			
Influencer FE	Y			
Month-Year FE	Y			
Game developer FE	Y			
Other Time FE	Y			

Table 3: OLS Regressions; standard errors in parenthesis clustered at influencer level. Influencer characteristics also include variables about most commonly played game. Game characteristics include genres, themes, and game modes.

Results of the regression are in columns 1 and 2 of Table 3. I also run a similar regression to Equation 3, except I subset just the sponsored game dev observations in columns 3 and 4. The regressions imply that game dev sponsors decrease ACV by $\approx 6\%$ and the number of followers acquired by $\approx 50\%$. Even though I do not control for selection into sponsored content, institutional detail supports the idea that the OLS measurements are some sort of upper bound. Streamers should always be trying to maximize ACV when making sponsored content because embedded ads and donations are positively correlated with ACV. Streamers are also in some sort of repeated game with potential sponsors; if a sponsor can see that streamers have shirked/sabotaged sponsored content efforts in the past, they would be hesitant to offer future deals to the streamer. Given these arguments, the treatment effect of an experiment where streamers were randomly assigned and forced to produce sponsored content would likely produce much more negative effects.

When streamers choose high disclosure, their ACV increases 12% over low disclosure and 6% over organic content. Focusing on the sample of only game dev sponsored streams, high disclosure is correlated with a still significant, but smaller 5% increase in ACV versus low disclosure.

I test for a partially separating equilibrium where θ_H types choose $j = HD$ and the other types choose LD. Note that ACV is essentially a measure of the product $P(engage|j)v(\theta)$ because both clicking into the stream ($k = engage$) and viewing time/intensity $v(\theta)$ affect ACV. In such an equilibrium, the coefficient β_a in the full sample is equivalent to the following:

$$\beta_a = P(engage|LD)v(\theta_L) - P(engage|LD)v(\theta_0)$$

I show $\beta_a < 0$, so $\theta_L < \theta_0$ since v is an increasing function.

The coefficient β_d is equivalent to:

$$\beta_d = P(engage|HD)v(\theta_H) - P(engage|LD)v(\theta_L)$$

This estimate is positive, but does not disentangle the effects of $P(engage|HD)$ and $v(\theta_H)$. Summing the two equations, we get:

$$\beta_a + \beta_d = P(engage|HD)v(\theta_H) - P(engage|LD)v(\theta_0)$$

The coefficient $\beta_a + \beta_d$ is positive in magnitude, but insignificant. In a world where $\theta_H \approx \theta_0$ or θ_H is slightly smaller than θ_0, this implies that $P(engage|HD) \geq P(engage|LD)$, one of the necessary conditions of a partially-separating equilibrium. As I will show in Section 4.3, this assumption on θ is not unreasonable in my setting. This finding, combined with the summary statistic that only 14.5% of game dev sponsored content is high disclosure (Table 1), are observable characteristics of a partially-separating equilibrium and provide a first step towards empirically validating the existence of the equilibrium.

4.2 Causal effects of disclosure and the single-crossing condition

In addition to the observed patterns on ACV and disclosure frequencies, the single-crossing condition must hold in a partially-separating equilibrium, i.e. the reputation costs of high disclosure must be higher for θ_L types than for θ_H types, $c(HD|\theta_L) > c(HD|\theta_H)$. I now provide an econometric test for the empirical existence of single-crossing in my setting.

Influencer reputation, like a sponsor's "alignment," is an unobservable object, so I use the change in the number of followers as a proxy for reputation costs of sponsorship and disclosure. Gaining followers is not a cost, so $c(\cdot)$ is negative follower change. For $j = LD$ streams, follower change corresponds to the $v(\cdot)c(\cdot)$ object because engagement has to happen in order for people to follow/unfollow. For $j = HD$ stream, we add c_{HD} because disclosure could spur unfollowing without engagement during the livestream. It is important to note that $P(engage|j)$ is irrelevant as Twitch follower change metrics are fully realized in the data and not forecasted. Follower metrics capture everyone who has followed or unfollowed regardless of if they engaged with the stream. The equilibrium dictates that only θ_H types choose high disclosure, so the OLS coefficient, β_d, captures only a function of $c(HD|\theta_H)$. The cost of high disclosure for low types, $c(HD|\theta_L)$, is completely unobservable in the OLS regression.

To address the unobservable nature of $c(HD|\theta_L)$, I use an IV regression. Intuitively, the estimate from an IV regression is the answer an experiment that asks, "what happens if all types were forced to disclose?" I instrument for the disclosure decision using the instrument: how often is a sponsored game disclosed by other influencers within the past thirty days. That is, for influencer i playing game g at day t, the instrument is:

$$z_{it} = \frac{\sum_{j \neq i} \sum_{\tau=t-30,...,t-1,t} HD_{j\tau} * ad_{j\tau} * 1\{g_{j\tau} = g_{it}\}}{\sum_{j \neq i} \sum_{\tau=t-30,...,t-1,t} ad_{j\tau} * 1\{g_{j\tau} = g_{it}\}} \tag{4}$$

The instrument tries to proxy a directive to disclose from a game developer by looking at the behavior of other influencers who may be in the same marketing campaign. Relevance for this instrument comes from the idea that the disclosure decision may be exogenously given by the developer for a particular campaign. Certain developers may be afraid of regulation enforcement or are just more likely to mandate disclosure for whatever reason. The more frequently other influencers have disclosed in the last thirty days, the more likely the focal influencer has also been told to disclose by the developer. The exclusion restriction comes from the idea that game developers are not leveraging disclosure as a lever that affects campaign performance. The same game would have different "alignment" across influencers, and developers simultaneously work with too many influencers to dictate which individuals signal and which do not signal. Therefore, some influencers are forced to prominently disclose a low type game which in the absence of the developer's instructions would have been diminished.

One threat to exclusion may be that marketing efforts can affect viewership outside of disclosure. I attempt to control for this threat by including a control that counts the number of instances a game is observed in the past week, the idea being that any other marketing efforts would be captured by a collective supply-side response from streamers to cash in on the temporary increase in profitability.

In column 4 of Table 4, I find that the effect of prominently disclosing game dev sponsors on follower change is significantly negative: $\beta_d^{IV} = -1.586$. Under very restrictive assumptions like homogeneous treatment effects (Blandhol et al., 2022), we can interpret β_d^{IV} as the overall effect of high disclosure on reputation. This interpretation means that the IV coefficient is some weighted average of $c(HD|\theta)$ over the distribution of θ in the population. Recall that both the theoretical model and the regression specifications restrict θ to be discrete with limited types. Then, only one

	Y: Log. Avg. conc. viewers		Y: IHS Follower Change		Y: High disclosure
	OLS	IV	OLS	IV	
Game dev spon hi. disc.	0.050	0.225	-0.026	-1.586	-
	(0.017)	(0.133)	(0.089)	(0.677)	-
Same-week streams	0.022	0.017	0.204	0.223	-
	(0.005)	(0.007)	(0.028)	(0.029)	-
Inst: % other disclose	-	-	-	-	0.220
	-	-	-	-	(0.027)
R^2	0.907	-	0.605	-	0.458
nobs	11989	11989	11989	11989	11989
Influencer Characteristics			Y		
Game Characteristics			Y		
Stream Characteristics			Y		
Influencer FE			Y		
Month-Year FE			Y		
Game developer FE			Y		
Other Time FE			Y		
First stage partial F:	-	-	-	-	66.0

Table 4: IV Regressions; standard errors in parenthesis clustered at influencer level. Influencer characteristics include number of followers and most frequently played game. Game characteristics include game age, genres, themes, and game modes. Stream characteristics include stream length, drops, tournament, championship, giveaway, charity, subathon, and first game of the day

additional assumption on $c(LD|\theta)$ is needed to make the following claim:

Proposition 1 *Under the discrete distributional assumption on θ, strict assumptions on the IV coefficient interpretation, and an assumption that $c(LD|\theta_H) = c(LD|\theta_L) = c$, a decreasing differences condition is satisfied if $\beta_d^{IV} > \beta_d^{OLS}$.*

See Appendix A.5 for the proof. Notably, the assumption on equivalent low disclosure costs has already been already made in the theoretical model (Section 3).

Comparing the third and fourth columns in Table 4 where the dependent variable of interest is follower change, I find that $\beta_d^{IV} < \beta_d^{OLS}$. Again, follower change is not a cost because an increase in the number of followers helps the influencer. Instead, losing followers is equivalent to a positive cost, so the correct comparison to make in my setting is $-\beta_d^{IV} > -\beta_d^{OLS}$. By Proposition 1, we have decreasing differences in $c(\cdot)$.

To statistically test for a difference, I employ a Hausman test that checks if the coefficient on the residuals from a "control function" regression is significantly different from zero because I only have one endogenous variable (Hansen, 2022). The alternative hypothesis of this test states that the 2SLS estimates are statistically different from the least squares estimates of the structural equation. To do this, I re-run the OLS regression (Equation 3), but include the residuals from the first stage of the IV regression as an additional x_{it}. The coefficient on the residuals has a coefficient 1.589 with a standard error of 0.648 and a t-statistic of 2.45.[19] I reject the null hypothesis of exogeneity at the 5% level, which means that the IV and OLS estimates are statistically different. This result is consistent with the theoretical model and supports our decision to treat high disclosure as an

[19]The same Hausman test with ACV as the dependent variable returns a t-stat of -1.34

endogenous variable. Decreasing differences in $c(HD|\theta)$ means that $c(HD|\theta_L) > c(HD|\theta_H)$, and the single-crossing condition for our utility function is satisfied.

4.2.1 Causal viewership effects of high disclosure

One result worth discussing is the positive coefficient of β_d^{IV} on ACV in the IV regression. This result is consistent with the theoretical model's prediction that high disclosure can increase viewership. Column 2 of Table 4 shows that the unbiased measurement of high disclosure on ACV is positive when considering all types. The point estimate is even larger than the OLS estimate in the column 1, which is measuring the ACV effect when only θ_H types are disclosing in equilibrium. In terms of ACV, β_d^{IV} is equal to the following:

$$\beta_d^{IV} = P(engage|HD)[p_H v(\theta_H) + p_L v(\theta_L)] - P(engage|LD)v(\theta_L)$$

Recall that $P(engage|HD) > P(engage|LD)$, and the convex combination $p_H v(\theta_H) + p_L v(\theta_L) > v(\theta_L)$. Provided that beliefs do not update in the IV, then the theoretical model would predict that high disclosure can increase ACV.

This positive measurement can be attributed to consumers' belief in equilibrium that only θ_H choose high disclosure. When off-path equilibrium beliefs have not updated, viewers are more likely to click the stream believing it to be θ_H. The IV regression exactly captures the off-path equilibrium outcome, as it recovers the true effect of high disclosure on ACV holding fixed all other features – like beliefs – in the environment. From the follower change analysis above, I even observe the punishment that viewers bring in the IV environment by not following/unfollowing streamers who select high disclosure for their θ_L alignment sponsor.

Another reason for this measurement could be an artifact of the setting. A common way to unfollow a streamer on Twitch would be to click into their livestream and unfollow from the page. Simply clicking into high disclosure streams would increase ACV, even if it is just to unfollow. This is a plausible explanation given the significant negative effect on follower change.

If we compare β_d^{IV} to β_d^{OLS}, the theoretical model predicts that $\beta_d^{IV} < \beta_d^{OLS}$ since $v(\theta_H)$ becomes a convex combination of $v(\theta_H)$ and $v(\theta_L)$ in the IV regression. This result does not actually hold in the data, as β_d^{IV} is larger than β_d^{OLS} in magnitude. However, the OLS coefficient has t-statistic almost twice as large and the two coefficients are not significantly different from each other.[20] Therefore, I note this discrepancy as a mere artifact of the setting and noise in the data. I remain confident in the theoretical model given the other empirical tests and findings.

4.3 Selection on unobserved alignment

I provide more data-driven evidence that streamers are not randomly disclosing; rather, they are selectively choosing disclosure levels based on sponsor type. The theoretical model treats sponsor type as a pure vertical quality, but that might not be the case in practice. "High type" sponsors

[20] See previous footnote

also depend on the horizontal match value - a kid-friendly streamer may not want to play any type of horror game no matter how good the vertical quality of the game/game developer might be.

I construct a proxy metric for this horizontal aspect of sponsor type or “alignment” by using qualitative characteristics of video games. The IGDB data comes with details about genres/themes/keywords of almost all video games seen on Twitch. I am able to compute, at every observational period, the prior frequency of genres and themes of games that a streamer has previously played. As an example, at the eleventh observation of a streamer, if the streamer has played platform games 9 out of the previous 10 observations I assign a value 0.9 to the platform genre.[21] There are 23 unique genres, 22 unique themes, and 6 unique game modes in the IGDB data, so at each observation a streamer's type is the 54-vector (including no genre, no theme, and no game mode) of prior frequencies. I can then compute the correlational coefficient between the 54-vector of genres for each game and the 54-vector of historical frequencies for each streamer to obtain a single number on the interval $[0, 1]$ representing the “alignment” between streamer and game at a specific time period.

Some streamers brand themselves as being “variety” streamers; these streamers build their followings by playing all sorts of wacky games. For such variety streamers, how similar current games are compared to games played previously may not be a great proxy for alignment since these streamers intentionally look for novelty. To address this concern, I subset streamers who have played less than 12 unique video games, which is the 25th percentile (see Table 1) of unique games played across all streamers in my data. Their revealed preference for just a few games strongly speaks to what they enjoy and what their audience expects from them. The metric is more suitable with my focal subset of streamers where good sponsor alignment is more precisely defined by historical preferences.

Figure 4 plots the cumulative distribution function of genre similarities by disclosure level. The CDF for high disclosure streams is close to the CDF for organic streams meaning that among this subset of streamers, the profile of games played during high disclosure streams is very similar to the profile of games played during organic streams. Most of the mass is towards the right, indicating that high disclosure sponsored games are quite similar to the historical composition of games played by streamers. The CDF for low disclosure is separated from either of the other two, so games played during low disclosure streams seem different from organically chosen games. The low disclosure CDF has much more mass to the left, implying that streamers try to hide dissimilar sponsors. I interpret this finding in two ways. First, $\theta_H \approx \theta_0$, especially in terms of horizontal similarity. Second, this is evidence for selection into high disclosure on brand alignment. Obfuscating poorly aligned games makes sense - streamers can hide poorly aligned sponsors to the fullest extent they are allowed to instead of obviously appearing as a “sellout.” When the sponsored game is better aligned, streamers are more willing to prominently disclose because the game is better suited to their expertise and/or their audience's taste.

[21] A game can be a part of multiple genres

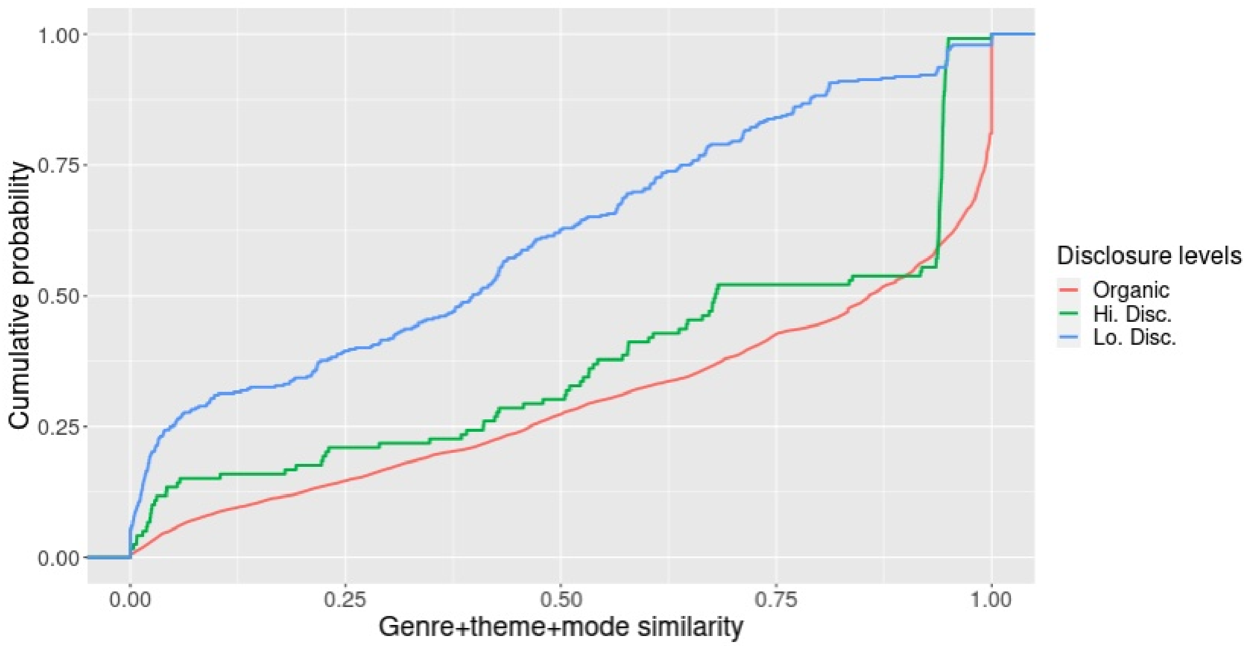

Figure 4: CDFs of genre+theme+mode similarity by disclosure level, streamers with ≤ 12 unique games played

4.4 Further evidence for selection: event studies

I isolate sponsorship events to analyze the context surrounding streamers' disclosure decisions using an stacked event-study framework. I use pre-trends to assess the selection into disclosure, and then use post-trends to assess the dynamic reputation costs $c(\cdot)$. The context around a sponsorship may provide further insight into why streamers choose high disclosure.

I aggregate stream-level data up to the week level to get a better sense of longer-term trends. I define "event" weeks where streamers have chosen either only low disclosure or only high disclosure for all of their sponsored game dev streams during that specific week. These events must be surrounded by a four-week radius where the streamer has no other sponsored game dev streams of any kind. For each disclosure type (high and low), week-by-week dynamic effects are recovered by running a regression of the form:

$$Y_{it} = \sum_{\tau=-4}^{4} \gamma_\tau 1\{t - E_i = \tau\} + \beta x_{it} + \alpha_i + \lambda_t + \varepsilon_{it} \quad (5)$$

where i is influencer, t is the week, and τ is the week relative to the sponsorship. E_i is the week when i performs the sponsorship. Y_{it} is the outcome of interest (e.g. ACV or follower change), x_{it} includes streamer characteristics like number of followers at the beginning of the week, number of hours streamed, and number of streams that week, α_i are streamer fixed effects, λ_t are week fixed effects, and ε_{it} is the error term. The γ_τ coefficients are the dynamic effects of the sponsorship event. I run this event study using dependent variables ACV and follower change, applying the imputation estimator from Borusyak et al. (2024).

If high and low disclosure decisions were randomly assigned at the time the sponsored content was broadcasted, we would expect no significant difference in pre-trends for both high and low disclosure events. Significantly differing pre-trends suggest that something leading up to the sponsorship week is affecting streamers' choices of sponsors and disclosure decisions. Post-sponsorship, if $c(HD|\theta) > c(LD|\theta)$ as the theoretical model suggests, random assignment of disclosure should create a significantly negative effect on follower change for high disclosure events. In addition, we should expect ACV to decrease in future periods as a result of longer term brand and reputation damage.

Figure 5a displays the stacked event study with ACV as the dependent variable.[22] I plot two trend lines, one for high and one for low disclosure events. In the week before the sponsorship event, $\tau = -1$, streamers who choose high disclosure command significantly higher ACV for their organic streams versus streamers who choose low disclosure. High disclosure streamers' 95% confidence interval does not include zero and does not overlap with low disclosure streamers' interval. It is unlikely that the sponsorship at $\tau = 0$ originated as an opportunity from the increased ACV at $\tau = -1$ because payment and contractual details take time to materialize. Additionally, many game dev deals revolve around game releases which are planned much farther out than one week in advance. Streamers must know something about the sponsorship that is preempting them to drive larger audiences the week before. Having a well aligned sponsor that you want to hype or show off to your audience could be a plausible explanation to try and increase viewership.

High disclosure streamers' increased ACV persists post-sponsorship event.[23] The 95% confidence intervals do not overlap one week and four weeks following the sponsorship event (and barely overlap at $\tau = 2$), so high disclosure streamers experience a significantly different ACV trend post-sponsorship. If θ_L types were the ones who chose high disclosure, or high disclosure was randomly assigned the theoretical model would predict a negative effect on ACV post-sponsorship as followers imposed reputation costs $c(HD|\theta_L)$ in a dynamic manner. The positive effect suggests that θ_H types are choosing HD.

To show that reputation does not contradict these results, I run the same event study but use weekly follower change as the dependent variable. I plot the results in Figure 5b. Crucially, none of the pre-trends or post-trends for high disclosure are significantly negative, and none of the trends are significantly different from each other.

Why might ACV matter for disclosure choice? One explanation could be that the industry contains dynamic relationships that are more complicated than what the theoretical model can capture. Streamers may be playing a static signaling game with their followers, but they can be engaged in a repeated game with sponsors, so obtaining higher ACV during a sponsored stream could lead to more lucrative sponsorship opportunities in the future. Streamers may exert more effort to drive higher ACV preceding a well-aligned sponsor, then using high disclosure to maintain ACV momentum and making themselves more attractive for future sponsors. I do not observe

[22] See Roth (2024) for discussion around visualizing the Borusyak et al. (2024) estimator.

[23] We do have to be careful with a causal interpretation given the significant pre-trend.

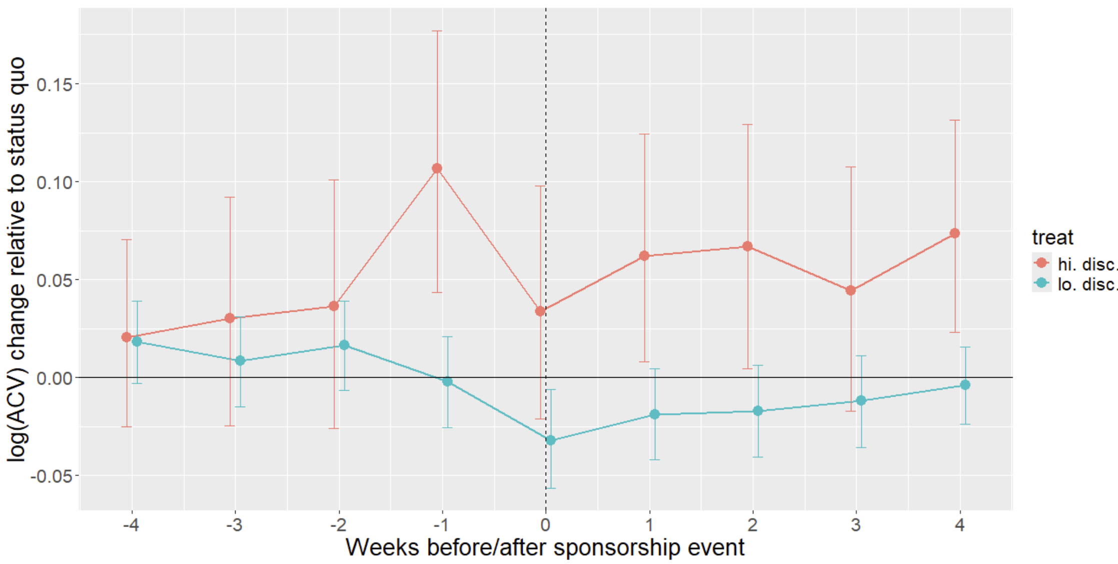

(a) ACV and disclosure choice. Low disclosure N: 1315. High disclosure N: 210.

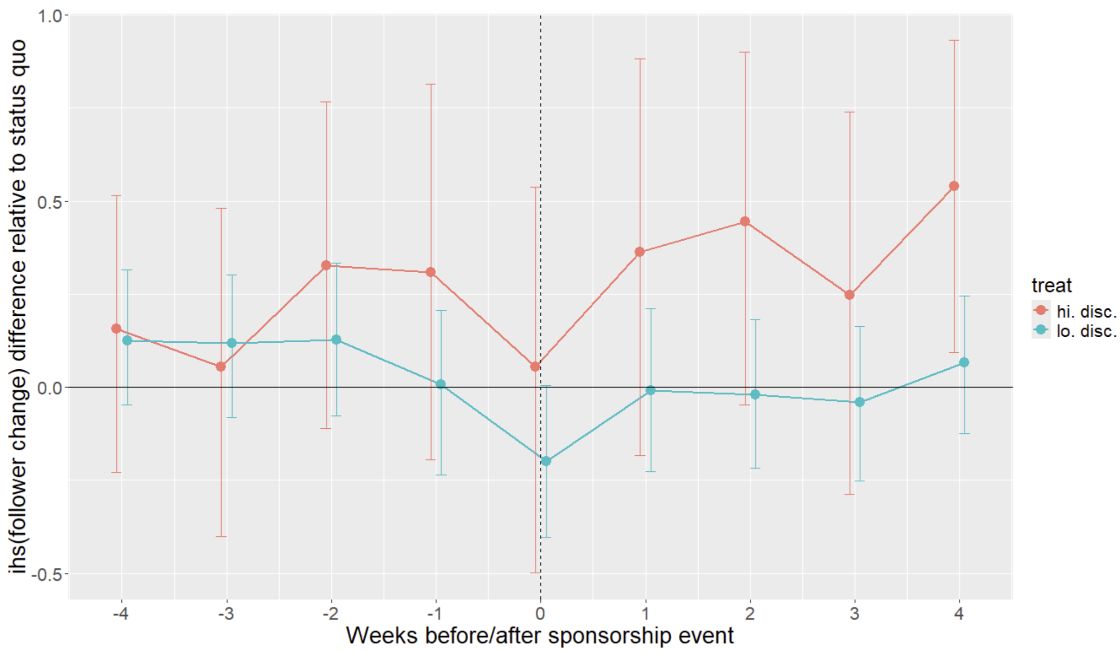

(b) Follower change and disclosure choice. Low disclosure N: 1315. High disclosure N: 210.

Figure 5: Event study – ACV and Follower change, grouped by disclosure choice. Coefficients estimated with Borusyak et al. (2024)'s imputation estimator. Error bars represent 95% confidence intervals around the coefficients and are computed using a Bayesian bootstrap of 1000 draws, with clustering at the streamer level.

sponsorship contracts, so capturing a dynamic relationship between advertisers and streamers is beyond the scope of this paper.

5 Discussion: Additional Generalizations

My theoretical signaling model can explain disclosure results from other marketing and economics contexts beyond Twitch. As discussed in Section 3.2, the model explains why TV shows and movies do not prominently disclose product placements and why radio DJs did not disclose payola.

Bairathi and Lambrecht (2023) find a similar effect in the context of Instagram influencers. Conditional on sponsorship, influencers who disclose their sponsored content have higher engagement rates than those who do not. The partially separating equilibrium in my model implies that disclosure is only observed because the Instagram influencer has high alignment with the sponsor. The influencer signals this high alignment to their followers by disclosing the sponsorship. Consequently, followers are more likely to engage with the content. The authors' measures of authenticity suggests that more authentic (i.e, better aligned) sponsorships mitigate the negative effect of sponsorship on engagement, suggesting that the single-crossing condition, $c(LD|\theta) > c(HD|\theta)$, may hold on Instagram. However, the authors do not have a measure of the cost of disclosure (e.g. follower change), so they do not test this directly.

Ershov and Mitchell (2023) find the opposite effect of disclosure, where post European regulation, disclosed sponsored content commanded less engagement than undisclosed sponsored content pre-regulation. The analysis conducted in Section 3.4 on regulating a low-disclosure pooling equilibrium supports this outcome: if $\theta_0 > \theta_H > \theta_L$, then engagement decreases post-regulation. This is possible in the early-mid 2010s because influencer marketing was still a nascent industry, so sponsored content that was as aligned as organic content may not have existed. Therefore, a low-disclosure pooling equilibrium existed in the pre-regulation period because payoffs for high disclosure were poor. In absence of regulation, sponsored content pooled with organic posts. Regulation forced influencers to choose high disclosure for sponsors, leading to the actions that perfectly revealed the θ_0 type. Engagement with sponsored content fell because consumers enjoyed organic content much more.

More generally, the model applies to digital platforms and search advertisements. Sahni and Nair (2020a) find, using a field experiment, that disclosure increases clickthrough and calls to advertising restaurants on a Yelp-like restaurant platform. They attribute this to a signaling effect for quality, whereby customers perceive advertising restaurants to be higher quality than non-advertising ones. My model suggests that "alignment" can also be an explanation. The advertising restaurant (sponsor) likely has high alignment with the Yelp-like platform (influencer) because the platform is known for being a good channel for consumers seeking to order food. So the platform obtains better engagement by disclosing the sponsorship to consumers. Had the sponsored search ad been something low alignment with the platform – say, something unrelated to food like a mechanic shop – then the model predicts that the platform would have incurred large reputation costs by disclosing this sponsorship. Consumers would believe that the platform is selling out by showing them the unrelated sponsored content. If the platform does not disclose the mechanic shop ad, the model still predicts a small hit to the platform (e.g. consumers thinking that a technical issues arose), but the platform can avoid the brunt of the brand equity costs.

6 Conclusion

This paper studies why voluntary disclosure of paid placements exists in some advertising contexts and not others. I introduce a signaling game where an influencer has a choice to disclose a paid placement or not, and a follower has a choice to engage with the content or not. I characterize the conditions of two types of equilibrium. I show that high reputation costs of high disclosure likely explain the existence of a non-disclosure pooling equilibrium in traditional product placement contexts like radio and TV. I find that a single-crossing condition on reputation costs, jointly with conditions on the probability of followers engaging with the content, are necessary for a partially-separating equilibrium where influencers with a high alignment sponsor choose high disclosure but those with low alignment sponsors choose low disclosure.

I collect over two years of online livestreaming data from Twitch.tv, recording all instances of prominent and diminished disclosure by popular Twitch streamers. Streamers voluntarily disclose some sponsorships from game developers but not others. I use these data to demonstrate the existence of a partially-separating equilibrium by reappropriating the Wu-Hausman test for endogeneity, comparing the coefficient of high disclosure on viewership and follower outcomes between OLS and IV regressions. Under certain assumptions the Wu-Hausman test empirically validates the single-crossing condition for the existence of a partially-separating equilibrium.

My theoretical framework can be extended to more general settings where paid placement occurs. It supports Coase's rationale that non-disclosure in radio is efficient because the disclosure imposes great costs on the listener and the radio DJ. My model may also explain why the disclosure of search ads may increase click through. Within the influencer marketing context, I can use alignment to explain observed patterns of disclosure on other prominent social media platforms like Instagram and Youtube.

The theory provides some limited insights into the discussion around the regulation and disclosure of paid placements. I show while regulating pooling equilibria may lead to decreased engagement in the short-run, regulating partially-separating equilibria increases engagement with organic content. My model falls short in assessing long-run outcomes of regulation because influencers cannot substitute between organic and sponsored content. A structural model of the influencer marketing industry that allows substitution between content types could provide more concrete answers regarding the topic of disclosure regulation.

References

Amano, T., Cheng, M., and Ofek, E. (2023). Interdependent creators of engagement: The role of co-followership patterns and content similarity in selecting influencers. *Available at SSRN 4319527*.

Aribarg, A. and Schwartz, E. M. (2020). Native advertising in online news: Trade-offs among clicks, brand recognition, and website trustworthiness. *Journal of Marketing Research*, 57(1):20–34.

Avery, J. and Israeli, A. (2020). Influencer marketing. *Influencer Marketing*.

Bairathi, M. and Lambrecht, A. (2023). Influencer marketing: Sponsorship disclosure and authenticity. *Available at SSRN 4679776*.

Berman, R. and Zheng, X. (2020). Marketing with shallow and prudent influencers. *The Wharton School Research Paper Forthcoming*.

Blandhol, C., Bonney, J., Mogstad, M., and Torgovitsky, A. (2022). When is tsls actually late? Technical report, National Bureau of Economic Research.

Boerman, S. C. (2020). The effects of the standardized instagram disclosure for micro-and meso-influencers. *Computers in Human Behavior*, 103:199–207.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, page rdae007.

Cheng, M. M. and Zhang, S. (2022). Reputation burning: Analyzing the impact of brand sponsorship on social influencers. *Shunyuan, Reputation Burning: Analyzing the Impact of Brand Sponsorship on Social Influencers (March 31, 2022)*.

Coase, R. H. (1979). Payola in radio and television broadcasting. *The Journal of Law and Economics*, 22(2):269–328.

Ershov, D., He, Y., and Seiler, S. (2023). How much influencer marketing is undisclosed? evidence from twitter.

Ershov, D. and Mitchell, M. (2023). The effects of advertising disclosure regulations on social media: Evidence from instagram. *The RAND Journal of Economics*.

Evans, N. J., Wojdynski, B. W., and Grubbs Hoy, M. (2019). How sponsorship transparency mitigates negative effects of advertising recognition. *International Journal of Advertising*, 38(3):364–382.

Fainmesser, I. P. and Galeotti, A. (2021). The market for online influence. *American Economic Journal: Microeconomics*, 13(4):332–72.

Hansen, B. (2022). *Econometrics*. Princeton University Press.

Hofstetter, R., Lanz, A., and Sahni, N. S. (2023). Contract design in influencer marketing.

Hristakeva, S. (2022). Vertical contracts with endogenous product selection: An empirical analysis of vendor allowance contracts. *Journal of Political Economy*, 130(12):3202–3252.

Karagür, Z., Becker, J.-M., Klein, K., and Edeling, A. (2022). How, why, and when disclosure type matters for influencer marketing. *International Journal of Research in Marketing*, 39(2):313–335.

Kay, S., Mulcahy, R., and Parkinson, J. (2020). When less is more: the impact of macro and micro social media influencers' disclosure. *Journal of Marketing Management*, 36(3-4):248–278.

Kim, D. Y. and Kim, H.-Y. (2021). Influencer advertising on social media: The multiple inference model on influencer-product congruence and sponsorship disclosure. *Journal of Business Research*, 130:405–415.

Li, N., Haviv, A., and Lovett, M. J. (2021). Digital marketing and intellectual property rights: Leveraging events and influencers. *Available at SSRN 3884038*.

Lin, Y., Yao, D., and Chen, X. (2021). Happiness begets money: Emotion and engagement in live streaming. *Journal of Marketing Research*, 58(3):417–438.

Lu, S., Yao, D., Chen, X., and Grewal, R. (2021). Do larger audiences generate greater revenues under pay what you want? evidence from a live streaming platform. *Marketing Science*, 40(5):964–984.

Mitchell, M. (2021). Free ad (vice): internet influencers and disclosure regulation. *The RAND Journal of Economics*, 52(1):3–21.

Morozov, I. and Huang, Y. (2021). Get rich on twitch: Economic analysis of video game streaming.

Nistor, C. and Selove, M. (2023). Influencers: The power of comments. *Available at SSRN 4118010*.

Nistor, C., Selove, M., and Villas-Boas, J. M. (2024). Influencers: To grow or to monetize. *Available at SSRN 4708107*.

Pei, A. and Mayzlin, D. (2022). Influencing social media influencers through affiliation. *Marketing Science*, 41(3):593–615.

Pöyry, E., Pelkonen, M., Naumanen, E., and Laaksonen, S.-M. (2021). A call for authenticity: Audience responses to social media influencer endorsements in strategic communication. In *Social Media Influencers in Strategic Communication*, pages 103–118. Routledge.

Rajaram, P. and Manchanda, P. (2020). Video influencers: Unboxing the mystique. *arXiv preprint arXiv:2012.12311*.

Roth, J. (2024). Interpreting event-studies from recent difference-in-differences methods. *arXiv preprint arXiv:2401.12309*.

Russell, C. A. (2002). Investigating the effectiveness of product placements in television shows: The role of modality and plot connection congruence on brand memory and attitude. *Journal of consumer research*, 29(3):306–318.

Sahni, N. S. and Nair, H. S. (2020a). Does advertising serve as a signal? evidence from a field experiment in mobile search. *The Review of Economic Studies*, 87(3):1529–1564.

Sahni, N. S. and Nair, H. S. (2020b). Sponsorship disclosure and consumer deception: Experimental evidence from native advertising in mobile search. *Marketing Science*, 39(1):5–32.

Schouten, A. P., Janssen, L., and Verspaget, M. (2020). Celebrity vs. influencer endorsements in advertising: the role of identification, credibility, and product-endorser fit. *International journal of advertising*, 39(2):258–281.

Simonov, A., Dubé, J.-P., Hitsch, G., and Rossi, P. (2020). State-dependent demand estimation with initial conditions correction. *Journal of Marketing Research*, 57(5):789–809.

Simonov, A., Ursu, R., and Zheng, C. (2021). Do suspense and surprise drive entertainment demand? evidence from twitch. tv. *Evidence from Twitch. tv (May 8, 2021). Columbia Business School Research Paper Forthcoming, NYU Stern School of Business Forthcoming.*

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374.

Sudhir, K. and Rao, V. R. (2006). Do slotting allowances enhance efficiency or hinder competition? *Journal of Marketing Research*, 43(2):137–155.

Sullivan, M. W. (1997). Slotting allowances and the market for new products. *The Journal of Law and Economics*, 40(2):461–494.

Yang, J., Zhang, J., and Zhang, Y. (2021). First law of motion: Influencer video advertising on tiktok. *Available at SSRN 3815124*.

A Appendix

A.1 Data Quality Issues

Since the scraped data comes from a public API, there are some inherent issues. The biggest issues both pertain to the measurement of the number of followers. Twitch struggles in distinguishing real human behavior from bot/automated computer behavior. Botting can be intrusive; bots can inflate the viewership and the number of followers of a channel to make it look more appealing to potential sponsors. Twitch sometimes conducts operations to delete bot accounts and remove bots from follower counts. Botting and bot-hunting can cause inaccurate, lumpy measures of followers.[24] I correct for potential "botted" data in the number of followers by identifying periods of bot-following and bot-deleting by Twitch using large jumps and dips, where the follower change is ± 5 standard deviations from a streamer's mean follower change, and construct a trend of "true" followers that a streamer has.

Another issue with the data involves channels that are not run by influencers. These channels often include the official channels of video game developers and publishers (e.g. Riot Games), dedicated esports tournament channels (e.g. ESL), and game-specific channels (e.g. Rainbow Six). I remove these Twitch channels because they rarely produced sponsored content and are not an individual brand. Twitch channels that are live for less than 10 days in the timeframe of the data are also removed.

Finally, there are some missing days in the data due to various issues with scraping. A few days are missing because of various server resets that the script was running on. A few days in August 2021 are missing because viewership data was bugged on the API endpoint for those days. A 30 day span from December 2021 - January 2022 is missing because the author went home for winter break and did not check if his scripts were still running. These create the following issues: for descriptive evidence, the biggest loss is sample size. However, over 95% of possible observations ($\sim 670,000$ observations) remain. For regressions, month-year fixed effects should handle any time-specific systematic biases. If video games streamed during these time periods are not fundamentally different from other time periods, other kinds of descriptive evidence should remain unbiased. For the dynamic model, one may worry about an "initial conditions" kind of problem occurring because of gaps in the data.[25] Biases related to this problem may be mitigated by the long panel; I observe over 700 days of choices for the initial cohort and about 600 days of choices for the additional cohort.[26]

[24] As an example: `https://twitchtracker.com/adinross`

[25] e.g. Simonov et al. (2020)

[26] Two additional features of the model directly address this. First, transitions are Markovian, mitigating effects on the transition likelihood. Conditional on observing today's state, yesterday's state only matters for computing the likelihood of observing such a transition. I can simply drop observations where I don't observe the prior day's state. Second, the model lacks any persistent components (including unobserved heterogeneity), so individual likelihoods do not have to be multiplied over time before taking logs.

A.2 Additional summary statistics

Statistic	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Avg. concurrent viewers	3,071.07	7,377.86	2.67	530.72	1,217.74	2,722.72	539,735.20
Stream length of game (hr) 4.57	5.67	0.50	1.92	3.67	6.17	714.75	
Followers gained	217.80	1,016.51	−118,777.00	3.00	30.00	140.00	283,152.00
Any ad indicator	0.044	0.205	0	0	0	0	1
Sponsored content indicator	0.018	0.133	0	0	0	0	1
Disclosed ad indicator	0.002	0.049	0	0	0	0	1

Table A.5: Observation level summary statistics

A.3 Pooling equilibrium proof

Lemma *At least one pooling equilibrium where $j = LD$ for every θ exists under the assumptions made in Section 3.1.*

The proof involves finding conditions under which the the incentive compatibility (IC) constraints of the influencers hold. I hold fix the follower's strategy, which is $k = engage$ if $\mathbb{E}[\theta] > g$ and $k = ignore$ otherwise. Because influencers do not know f's private cost of engaging, g, they must maximize their expected payoffs. The expected payoff IC constraints for influencers choosing $j = LD$ are:

$$P(engage|LD)v(\theta)\Big[1 - c(LD|\theta)\Big] > P(engage|HD)v(\theta)\Big[1 - c(HD|\theta)\Big] - (1 - P(engage|HD))c_{HD}$$

The probability of engaging depends on the posterior beliefs of f:

$$\begin{aligned} P(engage|LD) &= G(q(p_H\theta_H + p_L\theta_L) + (1-q)\theta_0) \\ P(engage|HD) &= G(\mu_{\theta_H,HD}\theta_H + \mu_{\theta_L,HD}\theta_L) \end{aligned}$$

I now show that IC constraints are easily satisfied under mild conditions. First consider the case when $P(engage|LD) \geq P(engage|HD)$:

$$\begin{aligned} P(engage|LD)v(\theta)\Big[1 - c(LD|\theta)\Big] &\geq P(engage|LD)v(\theta) - v(\theta)c(LD|\theta) \\ &> P(engage|LD)v(\theta) - v(\theta)\min_{\theta}\{c(HD|\theta)\} \\ &\geq P(engage|HD)v(\theta) - v(\theta)\min_{\theta}\{c(HD|\theta)\} \\ &\geq P(engage|HD)v(\theta)\Big[1 - c(HD|\theta)\Big] - (1 - P(engage|HD))c_{HD} \end{aligned}$$

Where the strict inequality in the second line comes from the monotonicity assumption on $c(\cdot)$. The last line holds true because we assumed $c_{HD} < \min_\theta v(\theta)c(HD|\theta)$. The proof is not conditional on a specific θ so IC constraints are satisfied for all θ types.

Below, we define conditions for when $P(engage|LD) \geq P(engage|HD)$. In this pooling equilibrium, no information is revealed by $j = LD$, so at the $j = LD$ information set $\mu_{\theta_H,LD} = qp_H$ and $\mu_{\theta_L,LD} = qp_L$. High disclosure $j = HD$ is off the equilibrium path so we must define beliefs at that node. Given monotonicity of CDF $G(\cdot)$, a necessary and sufficient condition for $P(engage|LD) \geq P(engage|HD)$ is:

$$\begin{aligned}
q(p_H\theta_H + p_L\theta_L) + (1-q)\theta_0 &\geq \mu_{\theta_H,HD}\theta_H + \mu_{\theta_L,HD}\theta_L \\
q(p_H\theta_H + (1-p_H)\theta_L) + (1-q)\theta_0 &\geq \mu_{\theta_H,HD}\theta_H + (1-\mu_{\theta_H,HD})\theta_L \\
qp_H(\theta_H - \theta_L) + (1-q)(\theta_0 - \theta_L) &\geq \mu_{\theta_H,HD}(\theta_H - \theta_L) \\
qp_H + (1-q)\frac{\theta_0 - \theta_L}{\theta_H - \theta_L} &\geq \mu_{\theta_H,HD}
\end{aligned}$$

The variables on the left hand side q, p_H, θ are all common knowledge. This scenario states that if off equilibrium path beliefs are made such that $f's$ expected payoffs are higher under $j = LD$ than $j = HD$, then influencers are better off pooling at $j = LD$.

Now let $P(engage|LD) \leq P(engage|HD)$. The pooling equilibrium could still exist because of large "sellout" costs associated with $j = HD$. If posteriors μ satisfy the inequality $c(HD|\theta) - c(LD|\theta) > \left[\frac{P(engage|HD)}{P(engage|LD)} - 1\right]$ for all θ, then:

$$\begin{aligned}
c(HD|\theta) - c(LD|\theta) &> \left[\frac{P(engage|HD)}{P(engage|LD)} - 1\right] \\
P(engage|LD)v(\theta)(c(HD|\theta) - c(LD|\theta)) &> (P(engage|HD) - P(engage|LD))v(\theta) \\
P(engage|HD)v(\theta)c(HD|\theta) - P(engage|LD)v(\theta)c(LD|\theta) &> (P(engage|HD) - P(engage|LD))v(\theta) \\
P(engage|LD)v(\theta)(1 - c(LD|\theta)) &> P(engage|HD)v(\theta)\Bigg[1 - c(HD|\theta)\Bigg] \\
P(engage|LD)v(\theta)(1 - c(LD|\theta)) &> P(engage|HD)v(\theta)\Bigg[1 - c(HD|\theta)\Bigg] - \\
&\quad (1 - P(engage|HD))c_{HD} \\
\pi(LD) &> \pi(HD)
\end{aligned}$$

We have found another sufficient condition for a $j = LD$ pooling equilibrium to exist. This conditions states that even if followers believe that their expected payoffs are higher under $j = HD$, influencers may still choose $j = LD$ if the costs of $j = HD$ are sufficiently high.

A.4 Partially-separating Equilibrium Proof

Lemma *At least one partially-separating equilibrium where $j = HD$ for θ_H and $j = LD$ for all other θ exists under the assumptions made in Section 3.1.*

Holding fixed the follower's strategy where $k = engage$ if $\mathbb{E}[\theta] > g$ and $k = ignore$ otherwise, the incentive compatibility constraints of θ_L types choosing $j = LD$ and θ_H types choosing HD are:

$$
\begin{aligned}
P(engage|LD)v(\theta_L)\Big[1 - c(LD|\theta_L)\Big] &> P(engage|HD)v(\theta_L)\Big[1 - c(HD|\theta_L)\Big] - \\
&\quad (1 - P(engage|HD))c_{HD} \\
P(engage|LD)v(\theta_H)\Big[1 - c(LD|\theta_H)\Big] &< P(engage|HD)v(\theta_H)\Big[1 - c(HD|\theta_H)\Big] - \\
&\quad (1 - P(engage|HD))c_{HD}
\end{aligned}
$$

In this partially-separating equilibrium, equilibrium actions require $\mu_{\theta_H,LD} = 0$ and $\mu_{\theta_H,HD} = 1$ because θ_H types only choose $j = HD$. The posterior beliefs of observing $j = LD$ are:

$$
\begin{aligned}
\mu_{\theta_L,LD} &= \frac{qp_L}{qp_L + (1-q)}, \\
\mu_{\theta_0,LD} &= 1 - \mu_{\theta_L,LD} = \frac{1-q}{qp_L + (1-q)}
\end{aligned}
$$

and the probabilities of engaging are:

$$
\begin{aligned}
P(engage|LD) &= G(\mu_{\theta_L,LD}\theta_L + (1 - \mu_{\theta_L,LD})\theta_0) \\
P(engage|HD) &= G(\theta_H)
\end{aligned}
$$

To simplify notation, let $P(engage|LD) = G_{LD}$ and $P(engage|HD) = G_{HD}$. Furthermore, recall that $c(LD|\theta_L) = c(LD|\theta_H) = c$ by assumption.

We define two conditions that together are sufficient for the equilibrium to occur. First, posterior beliefs μ and values of θ must exist in a manner such that $G_{HD} > G_{LD}$. The monotonicity of $G(\cdot)$ implies that q and p_L in the population must satisfy:

$$
\theta_H > \frac{qp_L}{qp_L + (1-q)}\theta_L + \frac{1-q}{qp_L + (1-q)}\theta_0
$$

The *single-crossing* condition is a sufficient condition for the existence of a partially-separating equilibrium. The single-crossing condition states that the reputation costs from choosing HD must be larger for θ_L types than for θ_H types. Formally, the single-crossing condition is:

$$
c(HD|\theta_H) < c(HD|\theta_L)
$$

To demonstrate this, we rewrite the IC constraints for both types to isolate G_{HD}. Consider the IC constraint for θ_L:

$$\begin{aligned}
G_{LD}v(\theta_L)[1-c] &> G_{HD}v(\theta_L)[1-c(HD|\theta_L)] - (1-G_{HD})c_{HD} \\
G_{LD}v(\theta_L)[1-c] &> G_{HD}(v(\theta_L)[1-c(HD|\theta_L)] + c_{HD}) - c_{HD} \\
G_{LD}v(\theta_L)[1-c] + c_{HD} &> G_{HD}(v(\theta_L)[1-c(HD|\theta_L)] + c_{HD}) \\
G_{HD} &< \frac{G_{LD}v(\theta_L)[1-c] + c_{HD}}{v(\theta_L)[1-c(HD|\theta_L)] + c_{HD}}
\end{aligned}$$

and the IC constraint for θ_H:

$$\begin{aligned}
G_{LD}v(\theta_H)[1-c] &< G_{HD}v(\theta_H)[1-c(HD|\theta_H)] - (1-G_{HD})c_{HD} \\
G_{LD}v(\theta_H)[1-c] &< G_{HD}(v(\theta_H)[1-c(HD|\theta_H)] + c_{HD}) - c_{HD} \\
G_{LD}v(\theta_H)[1-c] + c_{HD} &< G_{HD}(v(\theta_H)[1-c(HD|\theta_H)] + c_{HD}) \\
G_{HD} &> \frac{G_{LD}v(\theta_H)[1-c] + c_{HD}}{v(\theta_H)[1-c(HD|\theta_H)] + c_{HD}}
\end{aligned}$$

In this second inequality, note that because $c(HD|\theta_H) > c$, the right hand side fraction ignoring G_{LD} is greater than 1. Therefore, if $G_{LD} > G_{HD}$, the inequality could never hold, which is why $G_{HD} > G_{LD}$ is necessary for the partially-separating equilibrium to exist.

G_{HD} must be bounded by 1 above, which places a restriction on the maximum value of $c(HD|\theta_L)$:

$$\begin{aligned}
\frac{G_{LD}v(\theta_L)[1-c] + c_{HD}}{v(\theta_L)[1-c(HD|\theta_L)] + c_{HD}} &\leq 1 \\
G_{LD}v(\theta_L)[1-c] + c_{HD} &\leq v(\theta_L)[1-c(HD|\theta_L)] + c_{HD} \\
G_{LD}[1-c] &\leq [1-c(HD|\theta_L)] \\
c(HD|\theta_L) &\leq 1 - G_{LD}[1-c]
\end{aligned}$$

Piecing the two inequalities on G_{HD} together, we have:

$$\frac{G_{LD}v(\theta_H)[1-c] + c_{HD}}{v(\theta_H)[1-c(HD|\theta_H)] + c_{HD}} < \frac{G_{LD}v(\theta_L)[1-c] + c_{HD}}{v(\theta_L)[1-c(HD|\theta_L)] + c_{HD}} \tag{6}$$

Consider the following function and its partial derivatives:

$$\begin{aligned}
f(v,x) &= \frac{G_{LD}v[1-c] + c_{HD}}{v[1-x] + c_{HD}} \\
f_v(v,x) &= \frac{c_{HD}[G_{LD}(1-c) - (1-x)]}{(v[1-x] + c_{HD})^2} \\
f_x(v,x) &= \frac{v[G_{LD}v(1-c) + c_{HD}]}{(v[1-x] + c_{HD})^2}
\end{aligned}$$

Note that $f_v(v,x)$ evaluated at $x < 1 - G_{LD}[1-c]$ is negative. This means that for a fixed x,

$f(v(\theta_H), x) < f(v(\theta_L), x)$ since $v(\theta_H) > v(\theta_L)$.

Next, note that $f_x(v, x)$ is positive everywhere $v > 0$ and $x < 1$. So if $c(HD|\theta_L) > c(HD|\theta_H)$, then $f(v, c(HD|\theta_L)) > f(v, c(HD|\theta_H))$ for all v.

Combining these two together, we observe that $f(v(\theta_L), c(HD|\theta_L)) > f(v(\theta_H), c(HD|\theta_H))$, which is the requirement for both IC conditions to hold (Equation 6). We have already assumed that v is increasing. Hence, single-crossing, $c(HD|\theta_L) > c(HD|\theta_H)$, along with $G_{HD} > G_{LD}$ are sufficient for the partially-separating equilibrium to exist.

A.4.1 Alternative partially-separating equilibria

Here, I briefly discuss a partially-separating equilibrium where sponsored types perfectly separate from the organic type. That is θ_L and θ_H types both choose $j = HD$. Posterior beliefs at the $j = HD$ information set are $\mu_{\theta_L,HD} = p_L$ and $\mu_{\theta_H,HD} = p_H$. The probabilities of engaging are:

$$P(engage|LD) = G(\theta_0)$$
$$P(engage|HD) = G(p_H\theta_H + p_L\theta_L)$$

The IC constraint for both sponsored types is:

$$P(engage|LD)v(\theta)\Big[1 - c(LD|\theta)\Big] < P(engage|HD)v(\theta)\Big[1 - c(HD|\theta)\Big] - (1 - P(engage|HD))c_{HD}$$

If $P(engage|LD) > P(engage|HD)$, both IC constraints cannot possibly hold since $c(HD|\theta) > c(LD|\theta)$ by assumption and c_{HD} is positive, so the right hand has to be lesser. A sufficient condition for this is $\theta_H < \theta_0$.

In general, this equilibrium can exist if $\theta_H >> \theta_0 > \theta_L$, $p_L << p_H$, and $c(HD|\theta)$ is only a little larger than $c(LD|\theta)$. Under these circumstances, the probability of engaging when $j = HD$ is observed by the follower is much higher than when $j = LD$ is observed. This additional engagement outweighs the additional cost of high disclosure. In reality, this equilibrium is unlikely to exist. For example, if there existed a market where brand sponsors were incredibly desirable, many influencers would have already produced organic content related to such brands. There are very few brands that would prevent influencers from giving them free advertising and would want to maintain an image of exclusivity.

A.5 Decreasing differences proof

Proposition. *Under the discrete distributional assumption on θ, strict assumptions on the IV coefficient interpretation, and an assumption that $c(LD|\theta_H) = c(LD|\theta_L) = c$, a decreasing differences condition is satisfied if $\beta_d^{IV} > \beta_d^{OLS}$.*

Consider the specification:

$$c_{it} = \beta_0 + \beta_d HD_{it} + \beta_x x_{it} + \nu_i + \tau_t + \varepsilon_{it}$$

where c is the cost of disclosure and β_d is the coefficient on high disclosure. In our partially-separating equilibrium (Section 3.3), θ_H chooses HD, realizing costs $v(\theta_H)c(HD|\theta_H)$ if the follower engages and c_{HD} if the follower ignores. θ_L chooses LD, realizing costs $v(\theta_L)c(LD|\theta_L)$ if the follower engages.

Each i, t observation corresponds to one stream session. During a HD stream, the cost c includes the behavior from all followers who chose to engage with the stream session, but also all followers who ignored that stream session yet still observed the disclosure sponsored content. During a LD stream, the cost c only includes the behavior from all followers who chose to engage with the stream session, because those who did not engage did not discover the sponsored nature of the content.

Therefore in equilibrium, the measurement of c in the data corresponds to the sum $v(\theta_H)c(HD|\theta_H) + c_{HD}$ for HD choosing θ_H types and just $v(\theta_L)c(LD|\theta_L)$ for LD choosing θ_L types.

	θ_H	θ_L
Equilibrium action	HD	LD
Equilibrium reputation cost	$v(\theta_H)c(HD|\theta_H) + c_{HD}$	$v(\theta_L)c(LD|\theta_L)$

Table A.6: Equilibrium actions and reputation costs

If we estimate c using OLS, the interpretation of β_d^{OLS} is the difference in conditional expectation of c between $HD = 1$ and $HD = 0$. Equilibrium behavior of the two sponsored types allows me to rewrite this:

$$\beta_d^{OLS} = \mathbb{E}[\underbrace{v(\theta_H)c(HD|\theta_H) + c_{HD}}_{\theta_H \text{ eq. cost}} - \underbrace{v(\theta_L)c(LD|\theta_L)}_{\theta_L \text{ eq. cost}}] \tag{7}$$

If we estimate c_{it} using an instrumental variables regression where we instrument for HD_{it}, the interpretation of β_d^{IV}, under strict assumptions (homogeneous treatment effects, etc. see: Blandhol et al. (2022)), is the difference in unconditional expectations.

Using the assumed discrete distribution imposed on θ, we can rewrite this difference as a

Interpretation of β_d^{OLS}	$\mathbb{E}[c\|HD=1]-\mathbb{E}[c\|HD=0]$
	$=\mathbb{E}[\underbrace{v(\theta_H)c(HD\|\theta_H)+c_{HD}}_{\theta_H \text{ eq. cost}}-\underbrace{v(\theta_L)c(LD\|\theta_L)}_{\theta_L \text{ eq. cost}}]$
Interpretation of β_d^{IV}	$\mathbb{E}[c(HD)-c(LD)]$

Table A.7: Regression coefficient interpretations in a signaling equilibrium

weighted sum by the probability of θ in the population:

$$\beta_d^{IV} = \frac{p_H}{p_H+p_L}\mathbb{E}[v(\theta_H)c(HD|\theta_H)+c_{HD}-v(\theta_H)c(LD|\theta_H)]+ \\ \frac{p_L}{p_H+p_L}\mathbb{E}[v(\theta_L)c(HD|\theta_L)+c_{HD}-v(\theta_L)c(LD|\theta_L)]$$

where p_L and p_H are the probabilities of θ_L and θ_H in the population. Finally, we need the assumption:

Assumption 1 *Costs of low disclosure are similar for θ_H and θ_L types. That is, $c(LD|\theta_H) = c(LD|\theta_L)$*

This assumption is made in the setup of the signaling game and allows us to rewrite β_d^{IV} in an inequality as the convex combination of β_d^{OLS} and $\mathbb{E}[c(HD|\theta_L)-c(LD|\theta_L)]$:

$$\beta_d^{IV} = \frac{p_H}{p_H+p_L}\mathbb{E}[v(\theta_H)c(HD|\theta_H)+c_{HD}-v(\theta_H)c(LD|\theta_H)]+\frac{p_L}{p_H+p_L}\mathbb{E}[v(\theta_L)c(HD|\theta_L)+c_{HD}-v(\theta_L)c(LD|\theta_L)]$$
$$\beta_d^{IV} = \frac{p_H}{p_H+p_L}\mathbb{E}[v(\theta_H)c(HD|\theta_H)+c_{HD}-v(\theta_H)\underbrace{c(LD|\theta_L)}_{\text{assump.1}}]+\frac{p_L}{p_H+p_L}\mathbb{E}[v(\theta_L)c(HD|\theta_L)+c_{HD}-v(\theta_L)c(LD|\theta_L)]$$
$$\beta_d^{IV} < \frac{p_H}{p_H+p_L}\mathbb{E}[\underbrace{v(\theta_H)c(HD|\theta_H)+c_{HD}-v(\theta_L)c(LD|\theta_L)}_{\beta_d^{OLS}}]+\frac{p_L}{p_H+p_L}\mathbb{E}[v(\theta_L)c(HD|\theta_L)+c_{HD}-v(\theta_L)c(LD|\theta_L)]$$
$$\beta_d^{IV} < \frac{p_H}{p_H+p_L}\beta_d^{OLS}+\frac{p_L}{p_H+p_L}\mathbb{E}[v(\theta_L)c(HD|\theta_L)+c_{HD}-v(\theta_L)c(LD|\theta_L)] \tag{8}$$

Consider the case when $\beta_d^{IV} > \beta_d^{OLS}$. In this scenario, because the above inequality states that β_d^{IV} is less than a convex combination of β_d^{OLS} and $\mathbb{E}[v(\theta_L)c(HD|\theta_L)-v(\theta_L)c(LD|\theta_L)]$, β_d^{IV} must be less than $\mathbb{E}[v(\theta_L)c(HD|\theta_L)+c_{HD}-v(\theta_L)c(LD|\theta_L)]$. Using these inequalities on β_d^{IV}, we can derive the conditions for decreasing differences:

$$\mathbb{E}[v(\theta_L)c(HD|\theta_L)+c_{HD}-v(\theta_L)c(LD|\theta_L)] > \beta_d^{IV} > \beta_d^{OLS}$$
$$\mathbb{E}[v(\theta_L)c(HD|\theta_L)+c_{HD}-v(\theta_L)c(LD|\theta_L)] > \beta_d^{IV} > \mathbb{E}[v(\theta_H)c(HD|\theta_H)+c_{HD}-v(\theta_L)c(LD|\theta_L)]$$
$$\mathbb{E}[v(\theta_L)c(HD|\theta_L)-v(\theta_L)c(LD|\theta_L)] > \mathbb{E}[v(\theta_H)c(HD|\theta_H)-v(\theta_L)c(LD|\theta_L)]$$
$$\mathbb{E}[v(\theta_L)c(HD|\theta_L)-v(\theta_L)c(LD|\theta_L)] > \mathbb{E}[v(\theta_H)c(HD|\theta_H)-v(\theta_H)c(LD|\theta_H)]$$
$$\mathbb{E}[v(\theta_L)(c(HD|\theta_L)-c(LD|\theta_L))] > \mathbb{E}[v(\theta_H)(c(HD|\theta_H)-c(LD|\theta_H))]$$

The inequality holds between the third and fourth lines by Assumption 1 and because $v(\theta)$ is increasing and $\theta_H > \theta_L$. These properties of $v(\theta)$ also lead to the observation that the inequality in the final line implies

$\mathbb{E}[c(HD|\theta_L) - c(LD|\theta_L)] > \mathbb{E}[c(HD|\theta_H) - c(LD|\theta_H)]$. This is the condition for decreasing differences of $c(j|\theta)$

Ultimately, we are interested in comparing $c(HD|\theta_L)$ and $c(HD|\theta_H)$ to evaluate the single-crossing condition (Appendix A.4). Decreasing differences and Assumption 1 together imply $c(HD|\theta_L) > c(HD|\theta_H)$. In my setting, I use follower change as a proxy for reputation. Positive follower change is not a cost – it is a benefit to the influencer – so a positive cost $c(\cdot)$ requires *negative* follower change. To get decreasing differences, my setting necessitates that the IV coefficient β_d^{IV} is negative and significantly more negative than β_d^{OLS} since that implies a greater cost.

A.5.1 Equivalent low-disclosure cost discussion

Assumption 1: $c(LD|\theta_H) = c(LD|\theta_L)$ is discussed here. On Twitch, this assumption holds if those who unfollow/refuse to follow do so strictly because they are upset about the sponsored nature of the content. These could be users who click into the stream, read the full title and discover the sponsorship, and then immediately unfollow or leave. While not empirically verifiable with my data, the prominent positioning of the stream title underneath the stream in addition to a long right tail of watch times for streams in general makes this a plausible explanation. This assumption could also hold if on average, followers who respond more negatively to low-disclosure θ_L content are balanced out by followers who respond more negatively to high-disclosure θ_H content. I discuss both cases below.

Followers might be more willing to let low disclosure of θ_H types slide if they have watched for some time before discovering the sponsorship. Perhaps they enjoyed the content so they do not feel as upset as if they had discovered the sponsored nature earlier. In this case, $c(LD|\theta_H) < c(LD|\theta_L)$, so decreasing differences is harder to achieve because $c(HD|\theta_L)$ must be sufficiently larger than $c(HD|\theta_H)$ to offset the difference in $c(LD|\theta_H)$ and $c(LD|\theta_L)$.

On the contrary, followers may feel more upset when a θ_H type is obfuscating disclosure because they feel suckered into spending a long time viewing sponsored content. Or, they feel more likely to be deceived by subsequent θ_H content in the future. In this case, $c(LD|\theta_H) > c(LD|\theta_L)$, so decreasing differences is easier to achieve because $c(HD|\theta_L)$ does not need to be as large to offset the difference in $c(LD|\theta_H)$ and $c(LD|\theta_L)$.

It is unclear if which situation described above is more likely to be true for my setting. Because of this, I do not believe that the assumption $c(LD|\theta_H) = c(LD|\theta_L) = c$ is too restrictive and may be a fine approximation on average.

A.6 Robustness checks - descriptive evidence

Additional robustness checks for the descriptive evidence section (Section 4) are provided here.

A.6.1 High disclosure is immediately at the beginning of title

High disclosure is now defined as the hashtag occurring at the 0th character in the stream title. I run the OLS and IV regressions as before in the main text using the new definition of high disclosure. The results are shown in Tables A.8 and A.9.

	Full sample		Game dev sponsors only	
	Log ACV	IHS New Followers	Log ACV	IHS New Followers
Game dev sponsor	−0.066	−0.610		
	(0.015)	(0.080)		
Game dev spon hi. disc.	0.300	0.429	0.130	−0.130
	(0.045)	(0.158)	(0.034)	(0.156)
IHS game age	−0.022	−0.031	−0.011	−0.030
	(0.002)	(0.007)	(0.003)	(0.016)
Same-week streams	−0.003	0.074	0.021	0.205
	(0.003)	(0.012)	(0.005)	(0.028)
Log Stream Length	0.199	0.947	0.087	0.809
	(0.005)	(0.023)	(0.015)	(0.071)
Drops	0.223	0.833	0.115	1.097
	(0.019)	(0.069)	(0.032)	(0.147)
Tournament	−0.029	0.036	−0.017	0.888
	(0.015)	(0.042)	(0.041)	(0.243)
Championship	0.143	0.232	0.159	0.023
	(0.051)	(0.100)	(0.207)	(0.678)
Giveaway	0.053	0.294	0.073	0.584
	(0.014)	(0.054)	(0.026)	(0.135)
Charity	0.030	0.010	0.003	−0.425
	(0.021)	(0.078)	(0.054)	(0.335)
Subathon	0.100	0.032	0.294	0.603
	(0.034)	(0.082)	(0.069)	(0.268)
First Game	−0.223	0.096	−0.279	−0.277
	(0.009)	(0.027)	(0.020)	(0.220)
Log total followers	0.704	0.600	0.583	0.026
	(0.066)	(0.154)	(0.127)	(0.671)
Alignment	0.448	1.738	0.034	−0.031
	(0.035)	(0.140)	(0.058)	(0.332)
Num. obs.	668131	668131	11989	11989
R^2 (full model)	0.872	0.586	0.907	0.605
Game Characteristics		Y		
Influencer FE		Y		
Month-Year FE		Y		
Game developer FE		Y		
Other Time FE		Y		

Table A.8: OLS Regressions; standard errors in parenthesis clustered at influencer level. Influencer characteristics also include variables about most commonly played game. Game characteristics include genres, themes, and game modes.

The results mirror those in Tables 3 and 4 where high disclosure is defined with a looser cutoff at fifteen characters. The control function statistic for IHS New Followers is 1.290 with a standard error of 0.557, a

	Y: Log. Avg. conc. viewers		Y: IHS New Followers		Y: High disclosure
	OLS	IV	OLS	IV	
Game dev spon hi. disc.	0.130	0.192	−0.130	−1.351	-
	(0.034)	(0.113)	(0.156)	(0.568)	-
Same-week streams	0.019	0.018	0.165	0.179	-
	(0.006)	(0.006)	(0.032)	(0.032)	-
Inst: % other disclose	-	-	-	-	0.258
	-	-	-	-	(0.025)
R^2	0.907	-	0.605	-	0.518
nobs	11989	11989	11989	11989	11989
Influencer Characteristics			Y		
Game Characteristics			Y		
Stream Characteristics			Y		
Influencer FE			Y		
Month-Year FE			Y		
Game developer FE			Y		
Other Time FE			Y		
First stage partial F:	-	-	-	-	104.9

Table A.9: IV Regressions; standard errors in parenthesis clustered at influencer level. Influencer characteristics include number of followers and most frequently played game. Game characteristics include game age genres, themes, and game modes. Stream characteristics include stream length, drops, tournament, championship, giveaway, charity, subathon, and first game of the day

t-statistic of 2.32. The control function statistic for Log ACV is -0.065 with a standard error of 0.113, a t-statistic of -0.57. Both are consistent with the findings in the main text.