

Network Effects, Bargaining Power, and Product Review Bias: Theory and Evidence

Tom Hamami (Northwestern University)

28 December 2016

Abstract: I construct a theoretical framework for expert product reviews and demonstrate how the existence of positive network effects can make review inflation profitable even when fully rational consumers understand the existence of bias. This finding moreover suggests that product reviews, in addition to transmitting information, may also serve as a coordination mechanism for early adopters. Empirical application to video game review data suggests that this industry is in an inflation equilibrium. Specifically, I find evidence that reviews are inflated for games produced by large firms and for those that are part of pre-existing game franchises. Additionally, I find that review inflation is heterogeneous across genres that vary by the extent to which they produce network externalities, and I argue that this result is inconsistent with alternative explanations of review inflation.

Network Effects, Bargaining Power, and Product Review Bias: Theory and Evidence

Tom Hamami*

This version: October 18, 2016

Abstract

I construct a theoretical framework for expert product reviews and demonstrate how the existence of positive network effects can make review inflation profitable even when fully rational consumers understand the existence of bias. This finding moreover suggests that product reviews, in addition to transmitting information, may also serve as a coordination mechanism for early adopters. Empirical application to video game review data suggests that this industry is in an inflation equilibrium. Specifically, I find evidence that reviews are inflated for games produced by large firms and for those that are part of pre-existing game franchises. Additionally, I find that review inflation is heterogeneous across genres that vary by the extent to which they produce network externalities, and I argue that this result is inconsistent with alternative explanations of review inflation.

Keywords: Product reviews; Media bias; Regulatory capture; Video games.

JEL Classification Numbers: L14; L15; D21; D22; D62.

*Ph.D. Candidate, Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University. 2001 Sheridan Road, Evanston, IL 60208. Tel.: (410) 818-8472. Email: t-hamami@kellogg.northwestern.edu

“Do not patronize me by telling me the reader is the customer – your real customer is the one that pays you your revenue. And it is game industry advertisers.” – Trip Hawkins, former president of 3DO, in an email to *GamePro* in 2001.

1 Introduction

The idea of regulatory capture – that third party information disclosers, which ostensibly exist to serve the public, come under the thumb of those they are intended to regulate – dates back to the earliest literature on the economics of regulation.¹ Study of this phenomenon, however, focuses heavily on regulations of products and services that affect our physical and financial security. We do not often consider the notion of regulatory capture in media, but it seems almost self-evident that it occurs due to way the asymmetric relationship between information sources and journalists frequently makes it profitable for media outlets to report biased information. For example, when a National Football League player signs a new contract, initial media reports about the terms of the deal are often slanted favorably for the player. It is common knowledge among members of the media that the reason this occurs is because the initial leak comes from the player’s agent who wants to make himself look as successful as possible. Despite this common knowledge, every media outlet reports these terms – the sports media market, especially the market for breaking news, is highly competitive and no individual outlet can afford to forgo reporting this type of information.

This phenomenon is very hard to study because it is so difficult to quantify. There is one exception, however: expert product reviews often include numerical scores and follow this exact pattern – a producer creates a product and allows access to a reviewer who will then transmit a public signal about the product’s quality. As the marketing literature on product reviews suggests better reviews result in higher sales (see Basuroy and Ravid 2013,

¹See Stigler (1971); Peltzman (1976).

and Kaimann and Cox 2014, among many others), this implies the existence of motive and opportunity for producers to extort inflated reviews from reviewers. The reviewer requires access to the product to earn revenue and in many cases may also rely on advertising revenues from the producer. A producer therefore has the capacity for punishment (via blacklisting or withdrawal of advertising), and the conflict of interest becomes apparent.

More so than ever, product reviews play a critical role in consumer purchase decisions. Consumer surveys suggest that as many as 81% of consumers do online research prior to making a major purchase decision² and 88% trust [multiple] online reviews as much as a personal recommendation.³ In another survey, 90% of consumers recalled an instance where a positive review influenced their purchase decisions, and 86% recalled an instance where a negative review influenced their purchase decisions.⁴ Formal empirical research corroborates the findings of these surveys across many industries: in film (Basuroy et al. 2003; Elberse and Eliashberg 2003; Reinstein and Snyder 2005; Hennig-Thurau et al. 2006), finance (Goh and Ederington 1993), book publishing (Greco 1997), art (Caves 2000), wine (Hilger et al. 2011; Friberg and Gronqvist 2012), automobiles (Dewenter and Heimeshoff 2015), and video games (Zhu and Zhang 2010; Kaimann and Cox 2014). Since it is clear that expert reviews can influence demand and that the “journalism sub-industry” of many industries is set up in a way that promotes conflict of interest, research in this area should be of substantial interest to economists, marketers, and policy makers.

This article focuses on this conflict of interest. I construct a theoretical model of the review generation process where the producer has bargaining power over the reviewer. I first show that it is not obvious that review inflation makes a producer better off. This result follows from conventional economic reasoning: if consumers are rational, they understand and account for review bias in equilibrium. The conflict between this theoretical result

²Source: Survey and interviews conducted in 2013 by Rothstein Tauber, a marketing research firm.

³Source: 4th Annual Customer Review Survey conducted in 2014 by BrightLocal.

⁴Source: Survey conducted in 2013 by Dimensional Research.

and the empirical literature on expert reviews suggests the presence of another force at work. I subsequently show that, if a positive review actually creates value rather than merely transmitting information, the unique equilibrium involves maximal review inflation. Moreover, this result holds even if the value created is arbitrarily small. I then show that the presence of positive network externalities in some consumers' preferences can induce a coordination subgame where exactly this occurs. The key mechanic is that reviews can act as a coordination device for enthusiasts who must jointly decide which products to early adopt.

Although the coordination subgame in the model described above yields multiple equilibria, each type of equilibrium is observationally distinct. It should therefore be possible to examine the properties of a particular industry and determine in which coordination equilibrium it operates. As an empirical application of my theoretical results, I use review data to argue that the video game industry is in an inefficient inflation equilibrium. This inflation equilibrium predicts two key comparative statics. First, we expect more review inflation to occur for producers with more bargaining power. Second, we expect the degree of review inflation to increase with the prior probability that the product is of high quality. I use data on video game reviews to argue that both of these effects indeed hold in the video game journalism industry. Specifically, I compare expert reviews and user reviews on a sample of games. After controlling for several factors that correspond to biases in user reviews, I find that the difference between these two (i.e. the amount of expert review inflation) is greater on average for larger producers and for games from pre-existing franchises. Larger producers presumably have more bargaining power than smaller producers both in terms of total advertising spending and amount of content their products allow reviewers to produce, and I thus argue that this finding is consistent with the first comparative static. Moreover, since a sequel will only be made if its predecessor is sufficiently successful, and the quality of a sequel is positively correlated with the quality of its predecessor, we would expect that

the quality of a game that is a sequel is on average higher than one that is not. Thus, I argue that this result is consistent with the second comparative static. Finally, I show that inflated reviews in the video game industry are heterogeneous across genres that vary by the extent to which they produce network externalities and argue that this finding is inconsistent with alternative explanations of review inflation (such as consumer naivete).

These findings have interesting implications for consumer welfare in this market. In contrast with conventional economic wisdom, we observe a situation where increased competition in the market for reviews can actually decrease the quality of information transmitted. Moreover, the data suggest that the video game industry is in an equilibrium with review inflation; however, this equilibrium is Pareto dominated by another equilibrium where enthusiasts make their early adoption decisions independently of reviews. In fact, producers – the driving force behind the inflation – are actually strictly worse off than in an equilibrium with no inflation! This occurs because the positive network externalities in the model are a free lunch that is only being partially exploited in an inflation equilibrium, but the producer cannot unilaterally affect the coordination behavior of consumers. This result suggests the existence of heretofore unconsidered benefits of the “freemium” pricing structure that is growing in popularity in the software market. Under this structure, the producer of the software provides access to the application to consumers for free and subsequently earns revenues via “microtransactions” – selling cosmetic upgrades, additional content, and increased functionality. In the context of the current model, allowing consumers to learn the quality of a product at a reduced cost would maximize the extent to which the economy exploits these positive network effects.

This article is structured as follows: In the remainder of this section, I discuss related literature and how this article fits into existing scholarship on product reviews and persuasion. In Section 2, I introduce a generalized model of the review generation process and use three stylized examples to show how the results of the model can vary with the minor

modification of an assumption. I also discuss how the equilibrium selection of an industry can be inferred from observational data. In Section 3, I discuss further why the video game journalism industry is an ideal candidate for application of these results and use an empirical model to argue that this industry is in an inflation equilibrium consistent with the one described in the theory. In section 4, I show how heterogeneity of review inflation across genres is consistent with this theory and inconsistent with alternative theories of inflation. In Section 5, I discuss the welfare implications of these findings. I conclude in Section 6. Proofs of propositions can be found in Appendix A. Data details and robustness checks can be found in Appendix B. Discussions of theoretical robustness can be found in Appendix C.

Literature Review. This article is related to several strands of literature in the fields of economics and marketing. The theoretical framework contributes to a broad class of information transmission games first popularized by Crawford and Sobel (1982) and later revitalized by Gentzkow and Kamenica (2011). This literature is too vast to cover in its entirety; the model used in this article most nearly resembles that of Kolotinin (2015). In that model, the principal chooses ex ante the structure of the information transmission mechanism. The result is that the principal must trade off between the frequency of receiving a positive signal and the value of such a signal. This “quality vs. quantity” dynamic is precisely why, in (the first example of) the current model, review inflation does not necessarily benefit the principal when a review serves no purpose beyond information transmission. The current game also falls loosely into a subclass of games analyzed by Kartik (2009), where biased communication is costly and the equilibrium level of bias depends on the magnitude of this cost.

As mentioned above, a great deal of work has been done in marketing on product reviews – especially their impact on sales. A key finding in this literature is that the impact of expert reviews on sales persists despite the increased role played by physical and digital word-of-mouth in the dissemination of information (Basuroy and Ravid 2013). Cox and Kaimann

(2014) show evidence that positive expert reviews influence sales specifically in the video game industry.

A minor branch of this research area has dealt with detecting bias in expert reviews. Camara and Dupuis (2014) estimate a reputational cheap talk model using movie review data. Their key result – the finding of conservatism bias – is consistent with both the theoretical and empirical findings of the current article. The film industry, in general, has been popular for studying reviews due to the ubiquity of data (every local paper has a movie reviewer) and the econometrically convenient property that a movie ticket always costs the same. This industry is a poor choice to study the influence of producers on reviews, however; most movie reviews appear in general interest publications (such as newspapers) that do not rely excessively on the support of movie studios to survive. Hamami (2016) analyzes review data from a video game magazine owned by a game retailer and finds evidence of several forms of inflation. Further studies have found varying levels of bias in the book (Dobrescu et al. 2013), wine (Reuter 2009), automobile (Dewenter and Heimeshoff 2014), and fashion (Rinallo and Basuroy 2009) industries. The latter article is of particular interest; the authors study content selection rather than outright ratings bias, but they identify several factors that make an industry more susceptible to capture. I discuss these factors in the context of the video game journalism industry in Section 3.

Finally, this article relates to two broader literatures. The first is the study of general quality disclosure, whether disclosed by the producer or by a third party. Research in this area has focused primarily on regulations related to physical and financial security and less on consumer leisure industries. Dranove and Jin (2010) survey the existing state of theoretical and empirical literature on this topic. The second broader literature to which this article relates is the study of political slant in news media. Some research in this field focuses specifically on content and media firms' relationships with their advertisers. Ellman and Germano (2009) and Blasco et al. (2011) model two-sided markets where advertisers

suppress negative news about their products (for example, the negative health consequences of smoking tobacco).

The contribution of this article over existing literature is twofold. First, I introduce a theoretical framework that describes the review generation process and yields several insights about when and how review inflation benefits producers. Notably, I show that a positive review can actually create value by serving as a coordination device for enthusiasts who must jointly choose which products to early adopt; this finding provides an explanation for the empirical puzzle about the persistent impact of expert reviews on sales. Second, I evaluate the video game journalism industry for review bias. This industry in particular has been under increased media scrutiny over the past five to ten years, and I provide empirical evidence of what many have long suspected: that firms (either directly or indirectly) influence reviews of their products, and moreover, the manner in which this influence manifests is consistent with the network effects theory of inflation.

2 Model

2.1 General Model

There are three types of players: A production firm, a review firm, and a continuum of consumers with unit mass. The producer produces a product of random quality drawn from distribution G over support Q . The producer may contract with the reviewer to transmit information about the product's quality to consumers. The producer proposes an editorial policy⁵ $m : Q \rightarrow \Delta M$ to the reviewer, where M is a set of possible messages. Denote by \mathcal{M} the set of all possible editorial policies. If the reviewer declines this proposal, or if the producer declines to propose, we say $m = \emptyset$ and no information is transmitted to consumers.

⁵That the editorial policy is modeled as a binding contract is a simplification that allows us to capture an obviously repeated interaction in a single stage game. For more detailed discussion about dropping this assumption, see Appendix C.

If the reviewer accepts this proposal, message $m(q)$ is transmitted. Consumers, all of whom have unit demand for the product, observe message m and have the option to “early adopt” the product at early adoption price p chosen by the producer. Consumers who choose not to early adopt may wait and learn the product’s quality, at which point they have the option to purchase the product at wait price w , again chosen by the producer. The formal timing of the game is as follows:

1. The producer may choose to propose an editorial policy to the reviewer. This proposal is commonly observed. If the producer does not make a proposal, skip to stage 3.
2. The reviewer either accepts or declines the proposal of the producer.
3. Nature draws quality $q \in Q$ according to distribution G . No players immediately observe this draw, but its distribution is common knowledge.
4. If the two firms agreed on a proposal, $m(q)$ becomes common knowledge. If not, no information is transmitted.
5. Upon observing m (or not), the producer chooses a early adoption price $p \geq 0$. Consumers observe p and decide whether to early adopt the product at that price.
6. All players learn q . The producer chooses a wait price $w \geq 0$. Consumers who did not purchase in the previous stage observe w and decide whether to purchase at that price.
7. Payoffs are realized.

That the editorial policy is common knowledge is a direct implication of the fact that this one-shot specification is a reduced-form abstraction of a game that is played many times. Realistically, after the editorial policy is set, stages 3 through 7 are played repeatedly as the producer releases many products. Because consumers learn whether the reviewer lied at the end of each product cycle, even the beliefs of a consumer with a naive prior will converge to the correct policy if he is able to update those beliefs with Bayesian reasoning. The firms

therefore cannot systematically deceive consumers in this way. This assumption is common in the literature on media bias. For example, although news outlets mostly claim a lack of political bias in their reporting, the existence of political slant often becomes common knowledge through repeated interaction.

The producer has constant marginal production cost normalized to 0. The payoff for the producer is thus equal to its revenue as follows:

$$\pi = \gamma_p p + \gamma_w w \tag{1}$$

where γ_p and γ_w are the fractions of consumers who purchase at the early adoption and wait prices, respectively.

The payoff for the reviewer is zero if it is not contracted. If it is contracted, its payoff is:

$$v_r = R + c(G, m(\cdot)) \tag{2}$$

where R represents the lifetime discounted revenue stream earned by providing reviews of the producer's products and $c(G, m(\cdot))$ reflects expected lifetime reputational losses that depend on the average discrepancy between q and $m(q)$.

There exist up to two types of consumers. A fraction $\beta \in (0, 1]$ of consumers are “enthusiasts.” The remainder are “non-enthusiasts.” The payoff for any consumer who makes no purchase is zero. The payoff for consumers of type $i \in \{e, n\}$ who early adopt and for those who purchase at the wait price are, respectively:

$$v_p^i = \lambda_i u(q) - p \quad (3)$$

$$v_w^i = u(q) - w \quad (4)$$

Where $\lambda_i \geq 1$ and $\lambda_e > \lambda_n$. This reflects an “early adoption premium” such that, holding quality constant, consumers weakly prefer to early adopt than to wait, and the magnitude of this effect is larger for enthusiasts than for non-enthusiasts. Moreover, this premium can be constant or it might depend on other variables and parameters in the model. The extent to which this is the case will play a crucial role in the results derived below.⁶

Formally, a strategy for the producer consists of a proposal $m(\cdot) \in \mathcal{M}$, an early adoption pricing function $p : M \rightarrow \Delta\mathbb{R}_+$, and a wait pricing function $w : Q \times M \rightarrow \Delta\mathbb{R}_+$. Denote by \mathcal{P} and \mathcal{W} the sets of all possible pricing functions. A strategy for the reviewer is a function $r : \mathcal{M} \rightarrow \Delta\{\text{Accept}, \text{Reject}\}$ that either accepts or rejects any proposal made by the producer. A strategy for a consumer consists of two functions $b_p : M \times \mathbb{R}_+ \rightarrow \Delta\{\text{early adopt}, \text{wait}\}$ and $b_w : Q \times M \times \mathbb{R}_+ \rightarrow \Delta\{\text{purchase}, \text{no purchase}\}$ that determine whether he purchases the product at any stage of the game. The interim beliefs about the product’s quality for the producer and consumers are characterized by distribution H with density h over support Q . Denote $\mathcal{M}^{IR} \equiv \{m \in \mathcal{M} : R + c(G, m(\cdot)) \geq 0\}$. This is the set of all proposals for which it is weakly optimal for the reviewer to accept. Finally, let $\tilde{G}(\cdot)$ be the distribution over M induced by G and $m(\cdot)$.

Definition 1 *An equilibrium in this game consists of a producer strategy $(m^*(\cdot), p^*(m), w^*(q))$, a reviewer strategy $r^*(\cdot)$, consumer strategies $(b_p^{i*}(m, p), b_w^{i*}(m, w))$ for $i \in \{e, n\}$, and beliefs $h(q|m)$ such that for all $\hat{m}(\cdot) \in \mathcal{M}^{IR}$, $p \in \mathcal{P}$, and $w \in \mathcal{W}$, we have (arguments omitted):*

⁶In several places in this article, it might seem more logical for the premium to be additive rather than multiplicative. Modifying the assumption in this way does not affect the results.

$$\mathbb{E}_q[\gamma_p(m^*, p^*)p^* + \gamma_w(q, w^*)w^*] \geq \mathbb{E}_q[\gamma_p(\hat{m}(\cdot), p)p + \gamma_w(q, w)w] \quad [SR - P]$$

$$r^*(m(\cdot)) = \text{accept} \iff m(\cdot) \in \mathcal{M}^{IR} \quad [SR - R]$$

$$b_p^{i*}(m, p) = \text{early adopt} \iff \mathbb{E}_q[\lambda_i u(q) - p|m] \geq \max\{0, \mathbb{E}_q[u(q) - w^*(q)|m]\} \quad [SR - C_p]$$

$$b_w^{i*}(q, w) = \text{buy} \iff u(q) - w \geq 0 \quad [SR - C_w]$$

$$h(q|m) = \frac{\tilde{g}(m|q)}{\int^Q \tilde{g}(m|x)dx} \quad [Bayes]$$

2.2 Simplified Examples

We will now solve three simple versions of the general model. The examples are overly stylized by design in order to maximize clarity and intuition for the reader. I discuss the implications of relaxing several of the assumptions in this section in Appendix C.

In all settings, product quality will be binary (high quality or low quality), which will allow us to capture the extent to which the producer exerts its bargaining power with a single ordered pair. Further, in all settings, consumers have some variation of time preference (or else all consumers would simply wait for all uncertainty to resolve before making a purchase decision). It is in the nature of this time preference that the examples vary.

Formally, the following assumptions apply to all examples. $Q = M = \{H, L\}$, where H and L are positive numbers such that $H > L$. The prior probability that $q = H$ is $\theta \in (0, 1)$. The reviewer earns revenue R if it contracts with the producer. It suffers reputational losses

equal to $\phi > R$ if it transmits $m(q) \neq q$. We make the following assumption on R , θ , and ϕ .

Assumption 1 $\theta R + (1 - \theta)(R - \phi) < 0$

This assumption states that it is unprofitable for the reviewer to choose $m(q) = H$ for all q . This rules out a somewhat uninteresting case wherein the reviewer transmits no information but still earns revenues. We also assume that $u(q) = q$.

A key aspect of the parameter R is that it represents the marginal benefit of cooperation with the producer compared to the reviewer's outside option and thus captures the extent to which the producer has bargaining power over the reviewer. Although these parameters are exogenous, this framework is a lot more flexible than it looks. One can imagine more complicated machinery behind the exact relationship between the producer and the reviewer. For example, if the producer advertises directly through the reviewer, R can include the marginal difference in advertising revenue between the producer and the reviewer's next best option. R can also account for more complicated market structures. For example, if there are multiple producers whose products are substitutes in complicated ways, R simply represents the the difference in reviewer payoff in a partial equilibrium sense.

The bargaining process over editorial policy occurs as follows: The producer may choose to make a take-it-or-leave it offer $m : \{H, L\} \rightarrow \Delta\{H, L\}$. The reviewer can either accept or reject this offer. If the reviewer rejects, its payoff is zero and the remaining pricing and purchasing decisions are made on the basis of the prior alone.

Because quality is binary, an editorial policy m can be characterized by two numbers: the probability with which $m = H$ when $q = L$ and the probability with which $m = L$ when $q = H$. Denote these probabilities by η and ν , respectively. Throughout, I will refer to these numbers as the rates of review inflation and deflation, respectively. Thus, the bargaining process consists of the producer proposing an ordered pair (η, ν) and the reviewer accepting or rejecting this offer. It follows that:

- A producer strategy consists of $(\eta, \nu, p(H), p(L), p(\emptyset), w(H), w(L))$
- A reviewer strategy consists of accept/reject decisions for any $(\eta, \nu) \in [0, 1]^2$.
- A consumer strategy consists of purchase decisions for all possible early adoption prices (conditional on observing $m \in \{H, L, \emptyset\}$) and all possible wait prices (conditional on observing $q \in \{H, L\}$).
- Equilibrium interim beliefs about product quality are characterized as follows:

$$Pr(q = H|m = H) = \frac{\theta(1 - \nu)}{\theta(1 - \nu) + (1 - \theta)\eta} \quad (5)$$

$$Pr(q = H|m = L) = \frac{\theta\nu}{\theta\nu + (1 - \theta)(1 - \eta)} \quad (6)$$

2.2.1 Example 1

In this example, we assume that $\lambda_e > 1$ is an exogenous parameter and $\lambda_n = 1$. This is the simplest possible case; enthusiasts simply prefer to acquire the product sooner rather than later, and the size of the multiplier is independent of the product's quality and the reviewer's message.⁷ We will solve this game via backward induction. Consumers of either type who fail to early adopt will have willingness to pay in the wait stage equal to the ex post realization of q . Thus, in any equilibrium, we will have $w^* = q$ and consumers will purchase if and only if $w \leq q$.

The preceding implies that consumers anticipate that if they choose not to purchase in the early adoption stage, their payoff will be equal to zero. Therefore, non-enthusiasts will early adopt only if $\mathbb{E}[q|m] \geq p$ and enthusiasts will early adopt only if $\lambda_e \mathbb{E}[q|m] \geq p$. Since the producer is risk neutral, it is indifferent between selling to non-enthusiasts at $p = \mathbb{E}[q|m]$

⁷That non-enthusiasts have no time preference is an assumption made only for simplicity. In general, non-enthusiasts may have an early adoption premium so long as it is smaller than that of enthusiasts.

in the early adoption stage and at $w = q$ in the wait stage.⁸ However, the producer can sell to enthusiasts in the early adoption stage at a price as high as $\lambda_e \mathbb{E}[q|m] > \mathbb{E}[q|m]$. Thus, in any equilibrium, we will observe enthusiasts early adopting at price $p^* = \lambda_e \mathbb{E}[q|m]$ and non-enthusiasts purchasing at wait price $w^* = q$.

We know from Definition 1 that, in equilibrium, the reviewer accepts a proposal $m(\cdot)$ only if it satisfies its individual rationality constraint. When quality is binary, this constraint can be characterized as follows:

$$R - [\eta(1 - \theta) + \nu\theta]\phi \geq 0 \tag{7}$$

All that remains is to solve for the optimal editorial policy (η^*, ν^*) subject to the above constraint. This leads to the following result:

Proposition 1 *When $\lambda_e > \lambda_n = 1$ is a fixed parameter, producer profit is independent of η and ν . Thus, there exists an equilibrium with any $(\eta, \nu) \in \mathcal{M}$.⁹*

This result is driven by the fact that consumers understand the agreed-upon editorial policy. If the rate of review inflation is high, consumers respond by reducing their willingness-to-pay upon observing a good review. Therefore, even though review inflation yields a good review more often, the early adoption price the producer is able to charge is lower. When consumers are risk neutral, these two effects exactly offset. This is precisely the tradeoff between “quantity” and “quality” of information discussed by Kolotinin (2015).

How, then, might we go about predicting an outcome in this game? I offer two possible refinements, both of which leave truthful revelation (i.e. $\eta = \nu = 0$) as the sole surviving

⁸Recall that the review is as informative for the producer as it is for the consumer, so the producer’s expectation of the quality of the product updates from $\mathbb{E}[q]$ to $\mathbb{E}[q|m]$ upon observing the review. For the case where the producer possesses private information about the quality of the product, see Appendix C.

⁹Note that this does not violate Assumption 1. Equilibria with $(\eta, \nu) \in \mathcal{M} \setminus \mathcal{M}^{IR}$ involve the reviewer rejecting the producer’s offer.

equilibrium. The first is to adjust consumer preferences to include an infinitesimal amount of risk aversion.¹⁰ Functionally, this would result in consumers who behave as though they were risk neutral except in the case of indifference, when they choose the least risky option. In this case, enthusiasts strictly prefer waiting over early adopting at price $p = \mathbb{E}[q|m]$ except in the case of truthful revelation.

Alternatively, we can consider the quality of the relationship between the producer and the reviewer. Suppose that the working relationship between the two parties is strictly increasing in the payoff of the reviewer but that this factor is of second order importance to the producer.¹¹ Then, since the producer is otherwise indifferent between all sets of (η, ν) , the optimal proposal is again truthful revelation.

It is clear from this example that it is not obvious that the producer would benefit from review inflation. Moreover, there is a reasonable case to be made via refinements that the most likely outcome is one with no inflation at all.

2.2.2 Example 2

A key feature in the preceding example is that the review served no purpose except as a conduit for information. Suppose, instead, that a high review actually creates value in some way. Formally, we will capture this idea by assuming $\lambda_e(m = H) > \lambda_e(m = L) = \lambda_n(m = H) = \lambda_n(m = L) = 1$. It may not be immediately clear why such a situation would occur; Example 3 provides a setting in which this ordering arises endogenously.

In solving this version of the game, we can use the same logic as in Example 1 to establish that $w^* = q$ and that any consumer who has not yet made a purchase will do so in the wait stage. What occurs in the early adoption stage now depends on m . If $m = H$, then we are

¹⁰Formally, this could be modeled as $u(q) = q^{1-\frac{1}{k}}$, where k is an arbitrarily large positive number.

¹¹Formally, we could model this by expressing the producer's payoff as $\pi = \gamma_p p + \gamma_w w + \delta v_r$, where δ is an arbitrarily small positive number.

in the same situation as Example 1, and it must be that $p^* = \lambda_e(H)\mathbb{E}[q|m = H]$. If $m = L$, both the producer and consumers are indifferent between early adopting at price $\mathbb{E}[q|m = L]$ and waiting. Either can occur in equilibrium, but payoffs are the same in either case. We are again left only with deriving the optimal editorial policy subject to the reviewer’s individual rationality constraint.

Proposition 2 *When $\lambda_e(H) > \lambda_e(L) = \lambda_n(H) = \lambda_n(L) = 1$, the (essentially) unique equilibrium features maximal review inflation. Specifically, $\eta^* = \frac{R}{(1-\theta)\phi} > 0$ and $\nu^* = 0$.*

Assumption 1 ensures that the equilibrium value of η^* is strictly between 0 and 1. While the key assumption in this example may appear somewhat *ad hoc* (to be remedied in Example 3), it illustrates an important insight: When a positive review actually creates value rather than merely transmitting information, we expect producers to induce as much review inflation as their bargaining muscles allow. Importantly, this is true even though the fraction of consumers who early adopt (β) can be arbitrarily close to zero and the magnitude of the value created ($\lambda_e(H)$) can be arbitrarily close to 1. The equilibrium is essentially unique up to variations in early adoption pricing when $m = L$. Specifically, the producer can choose any $p^* \geq L$ since it is indifferent between selling at $p = \mathbb{E}[q|m = L] = L$ in the early adoption phase and $w = L$ in the wait phase. Notably, this equilibrium is robust to both the “risk aversion” and “relationship quality” refinements described in Example 1.¹²

In addition, this result provides two testable comparative statics that will play an important role in the empirical analysis. Specifically, the rate of review inflation:

1. Increases with the amount of revenue lost by the reviewer from not contracting with the producer. This parameter R directly captures how much bargaining power the producer possesses.

¹²In general, this equilibrium can survive a modest amount of consumer risk aversion. However, if risk aversion is sufficiently severe, the destruction of value from creating a risky asset dwarfs the creation of value from $\lambda_e(H)$, and the only equilibrium is truthful revelation.

2. Increases with the prior probability that the product will be of high quality. The higher this probability, the less likely it is that the reviewer will actually have to pay the reputational costs associated with lying.

2.3 Example 3

The preceding examples demonstrate that review inflation may benefit the producer only if a positive review actually creates value rather than merely transmitting information. In this example, we will see that the presence of positive network externalities can generate an equilibrium identical to the one from Example 2. Suppose now that λ_e depends on γ_p , the fraction of consumers who early adopt. Specifically, suppose we have $\lambda_e(0) = 1$ and $\lambda_e'(\cdot) > 0$ so that enthusiasts value early adoption more when a larger fraction of the population also early adopts. We can think of these network effects as literal – in the case of video games, many have multiplayer features such that the product literally has more value the more people play – or as a social effect: “Everyone I know owns this product, so I must own it, too.” As in the previous examples, we let $\lambda_n = 1$. Moreover, as in example 2, we again allow β to be arbitrarily close to zero and $\lambda_e(\gamma_p)$ to be arbitrarily close to 1. This means that two people discussing a product on an internet message board and enjoying that experience is sufficient to satisfy this assumption.

In addition, we impose the following restriction on consumer purchasing behavior:

Assumption 2 *Consumer purchasing strategies are characterized by cutoff (\bar{p}, \bar{w}) such that they are willing to purchase if and only if $p \leq \bar{p}$ (in the early adoption stage) or $w \leq \bar{w}$ (in the wait stage).*

This assumption is fairly innocuous, as this is the type of behavior we typically expect to see from consumers. As in the previous two examples, it is immediate that, in equilibrium, we have $w^* = q$ and all remaining consumers purchase at this price, so Assumption 2 has no bite

in the wait stage. In the early adoption stage, however, this will rule out bizarre equilibria where enthusiasts early adopt at price p but not at price $p - \epsilon$. Regardless, the optimal behavior of an individual enthusiast now depends on the behavior of other consumers, which in turn affects the optimal editorial policy.

Additionally, it is a natural assumption to allow both the revenue the reviewer can earn and the size of its potential reputational costs to depend on the magnitude of the equilibrium network externality. We will thus assume henceforth that R and ϕ are functions that depend on λ_e , though this assumption will not come into play until much later.

Proposition 3 *When $\lambda_e(0) = 1$, $\lambda'_e(\cdot) > 0$, and $\lambda_n = 1$, this induces a coordination subgame for enthusiasts that can result in five different types of equilibrium outcomes:*

1. *Enthusiasts early adopt if $m \in \{H, L\}$.
 π is independent of (η, ν) .*
2. *Enthusiasts do not early adopt for any $m \in \{H, L, \emptyset\}$.
 π is independent of (η, ν) .*
3. *Enthusiasts early adopt if $m = H$ but not if $m = L$.
 $\eta^* = \frac{R}{(1-\theta)\phi} > 0$ and $\nu^* = 0$.*
4. *Enthusiasts early adopt if $m = L$ but not if $m = H$.
 $\eta^* = 0$ and $\nu^* = \frac{R}{\theta\phi} > 0$.*
5. *Enthusiasts early adopt if and only if $m = \emptyset$.
 $m^*(\cdot) = \emptyset$.*

In the first two types of equilibria, enthusiast behavior is independent of m . This gives us a result identical to that of Example 1 (including truthful revelation being the unique survivor of our refinements). In equilibrium 3, however, enthusiasts early adopt only after observing a positive review. Thus, the benefit from the positive network externalities is only

realized in the case of a positive review. This gives us a result with maximal review inflation identical to that of Example 2 (and the associated comparative statics). Analogously, in equilibrium 4, enthusiasts only early adopt after observing a negative review, resulting in maximal review deflation. In equilibrium 5, enthusiasts early adopt only in the case of no review, predictably resulting in no reviews.

While equilibrium multiplicity might be an issue for the purposes of prediction, it is not a problem for evaluating existing industries because each equilibrium type is observationally distinct. This is useful information because equilibrium type 1 with a truthful editorial policy Pareto dominates all other equilibria by maximizing the positive network effects and minimizing reputational costs. Thus, observing an industry in any other equilibrium suggests the possibility that welfare could be improved through policy. The equilibria differ according to the following criteria, listed below and summarized in Table 1:

Table 1: Equilibria Properties

Eqm.	Pre-orders	Review Bias	Price Dispersion	Comparative Statics
(1)	Many	Little to none	High	None
(2)	Few	Little to none	High	None
(3)	Review Dependent	Inflation	Low	$\frac{\partial \eta}{\partial R}, \frac{\partial \eta}{\partial \theta} > 0$
(4)	Review Dependent	Deflation	Low	$\frac{\partial \nu}{\partial R} > 0, \frac{\partial \nu}{\partial \theta} < 0$
(5)	Many	N/A	N/A	None

- **Early Adoption Frequency:** We can use the prevalence of pre-orders as a proxy for the rate of early adoptions. In equilibria 1 and 5, we should observe large numbers of pre-orders irrespective of the the content of the review. In equilibria 3 and 4, we should observe many pre-orders only following a positive and negative review, respectively. In equilibrium 2, we should observe few pre-orders under any circumstance.
- **Review Inflation:** In equilibria 1 and 2, we should observe little to no review inflation. In equilibrium 3, we should observe high inflation. In equilibrium 4, we should observe high deflation. In equilibrium 5, there are no reviews at all. Equilibria 3 and 4 also

have unique comparative statics related to their rates of inflation/deflation that can be tested empirically.

- **Initial Price Dispersion:** In equilibria 1 and 2, early adoption price should be very different depending on the content of the review. This is because reviews are largely reliable in equilibrium 1, so producers can charge a high price after a positive review. In equilibria 3 and 4, price should vary relatively little based on the content of the review since consumers are not sure whether the review is truthful.

For example, we can examine the automobile industry in the United States and note that it fits the profile of equilibrium type 2. Pre-ordering is uncommon for most types of automobiles irrespective of quality. It is likely that there is relatively little review inflation,^{13,14} and initial prices align closely with quality.

In the next section, I will use review data to argue that the video game industry is in equilibrium 3, to which I will hereafter refer as an “inflation equilibrium.” While the presence of review inflation in the video game industry is itself interesting, these results also yield a new insight about the function of expert reviews: In addition to transmitting information about product quality, they also serve as a coordination mechanism for enthusiasts who must jointly decide which products to early adopt. This theory helps explain why the significant impact of expert reviews on revenues persists despite the increasing relevance and proliferation of user word-of-mouth via the internet.

¹³*Consumer Reports*, the most well-known source of expert reviews of automobiles, accepts no advertising and makes all purchases anonymously, costly sacrifices to signal its editorial independence. It thus has no firm-level motivation to bias reviews, though agent-level corruption could still exist.

¹⁴Dewenter and Heimeshoff (2015), mentioned in the introduction, specifically examines the German market.

3 Empirical Analysis

3.1 Why Video Games?

In this section, I discuss the reasons for choosing the video game industry as the empirical application of the theory presented above. Pre-orders of video games are common, and volume is affected by reviews.¹⁵ Price dispersion for newly released games is minimal; for example, the price of a new Xbox 360 game upon release in the United States was virtually always between \$50 and \$60.

Moreover, the video game industry is both representative and economically significant. According to the Entertainment Software Association 2016 Annual Report, 63% of households include at least one person who plays video games regularly. 41% of game players are female, and 46% are 36 years of age or older. The ESA estimates the total 2015 consumer spend on the video game software and hardware in the U.S. to be \$23.5 billion.

The next most obvious reason for choosing this industry is availability of data. Quantitative review data on video games is plentiful. In judging review inflation, I compare a game's Metacritic score (a score that aggregates a large number of expert reviews) with its Amazon user rating. The median game in my sample has a Metacritic score based on 44.5 expert reviews and has 60 Amazon ratings. In terms of sheer data, the video game industry is outdone perhaps only by the film industry. However, because film reviews are mostly published in general interest publications (such as newspapers, on whom movie studios presumably have little influence), the film industry is a poor choice for evaluating this particular dynamic.

Convenience, however, is not the only reason for choosing the video game industry. There exists a non-trivial number of claims of review bias related to this industry, and unrest has

¹⁵Cox and Kaimann (2014) demonstrate the impact of expert reviews on video game sales; Hernandez and Vicdan (2014) provide a qualitative discussion of motivations for pre-ordering video games, including informational sources such as reviews.

been growing among video game fans over the past five to ten years as a result. Moreover, Wingfield (2007) describes how an increasing number of game publishers tie developer compensation and/or job security to review scores. In this section, I will describe several anecdotes that serve as probable cause to investigate video game journalism.

On November 28, 2007, Jeff Gerstmann was terminated from his job as editorial director for the popular gaming website Gamespot. Immediately, rumors abound that this termination was a direct result of a poor review for recently released game *Kane & Lynch*. Eidos Interactive, the publisher of the game, had recently poured large amounts of advertising money into Gamespot, and many people in the industry and the public believed that Eidos pressured Gamespot into firing Gerstmann. On March 15, 2012, after the non-disparagement agreement with the company that owned Gamespot was nullified, Gerstmann confirmed in an interview that it was not only the *Kane & Lynch* incident, but a series of similar incidents that involved him being “called into the room to discuss a review” that resulted in his termination.¹⁶

In 2001, the game *Portal Runner* was given a poor review by the (now defunct) video game magazine *GamePro*. Trip Hawkins, the president of (now defunct) 3DO, the company that produced the game, wrote in an email that 3DO would cut back advertising in *GamePro* as a direct result of this review. Perhaps coincidentally, Hawkins was also the founder of Electronic Arts, which is now the largest video game software publisher in the world and whose games are particularly well received by expert critics.

Advertising is not the only way in which game producers can exert leverage over reviewers. After all, not all producers advertise with every reviewer. Reviewers also need the cooperation of producers to get early access to products so that they may create content for

¹⁶This example, along with many of the anecdotes from the video game journalism industry I discuss in this article, are described in greater detail in a *Forbes* article published in October 2012 by Erik Kain. Available at <http://www.forbes.com/sites/erikkain/2012/10/26/all-the-pretty-doritos-how-video-game-journalism-went-off-the-rails/>, accessed February 27, 2015.

readers. As a punishment for “undesirable behavior,” a producer can “blacklist” a reviewer by reducing or eliminating access to products.

In 2007, Kotaku, a popular website devoted to video games, published rumors about Sony’s upcoming Playstation 3 console against Sony’s will. As a response, Sony cancelled all executive interviews with Kotaku and dis-invited Kotaku from one of its media events. Kotaku got in hot water again in 2010 when it published an article alleging a corrupt relationship between Konami, a major Japanese game publisher, and *Famitsu*, the largest video game magazine in Japan.¹⁷ Konami subsequently revoked Kotaku’s invitation to its launch event for the game in question.

Metacritic is a website that aggregates expert reviews into a single numerical score. Metacritic has a strict policy that, when aggregating reviews, it always uses the first published score for each review. Metacritic claims that the reason for this policy is that it has become aware of instances where producers pressure reviewers to modify their initial reviews. Thus, using the first published review both maximizes the accuracy of the Metacritic score as a measure of expert sentiment and also mitigates the incentive producers have to pressure reviewers.

These anecdotes have not gone unnoticed by the gaming community. Tensions in this community boiled over in 2014 with the emergence of the ill-defined “GamerGate” movement. While the news media focused on the more politically explosive aspects of the controversy, at least some who rallied behind the GamerGate banner did so to raise awareness of what was at least *perceived* corruption in the video game journalism industry. This perception makes the industry worthy of investigation, irrespective of the actual result.

Finally, it is worthwhile to discuss the industry in the context of a study done by Rinallo and Basuroy (2009). The authors studied the relationship between content choice and

¹⁷Source: <http://www.1up.com/news/konami-shuns-blog-metal-gear>, accessed February 27, 2015.

advertising in the fashion journalism industry. They found that (a) there is a significant effect between quantity of coverage and advertising, (b) this effect increases with the extent to which a publication depends on a specific industry for advertising revenue, (c) the effect is more pronounced for larger and more innovative firms, and (d) these effects occur both in Europe and North America. The key finding for our purposes is (b); video game reviews appear overwhelmingly in publications and websites devoted specifically to video games. We would therefore conjecture that these outlets are more susceptible to producer influence than, say, a publication devoted to the entire entertainment industry (including video games). Result (c) is consistent with one of the key results from our theory and from the empirical analysis (described in the next subsection). Additionally, the current analysis similarly finds no distinction between American and European producers (see Appendix B).

3.2 Empirical Design and Results

In the previous section, I argued that there are two comparative statics that uniquely appear in an inflation equilibrium. Specifically, we should expect inflation to be higher for firms with high bargaining power and for products with higher prior probability of being high quality. In this section, I will use review data on 550 Xbox 360 games released between 2007 and 2013 to show that these comparative statics hold in the video game industry and its corresponding journalism sub-industry. The existence of review inflation is itself of interest, as the ensuing analysis appears to confirm suspicions long held by those in and around the industry. Moreover, the consistency of the data with the comparative statics generated by the theoretical model lends support to the idea that the impact of expert reviews on revenues is at least partially driven by their role as a coordination device for early adopters with network effects. For a detailed description of the data and a variety of robustness checks, see Appendix B.

To analyze review inflation, we must first define a way to measure it. To do so, we will compare expert reviews of video games against user reviews of video games. As a measure of how experts feel about a game, we will use the game’s Metacritic score. Metacritic is a website that aggregates expert reviews for movies, video games, and television programs into a single number between 0 and 100. This score has the benefit of capturing overall expert sentiment about a game while ironing out the idiosyncracies of individual expert reviewers. As a measure of how non-expert users feel about a game, we will use the games average rating on Amazon (normalized to the same 0-100 scale). User ratings of this sort are subject to their own biases, but producer influence is presumably not one of them. If we successfully control for these biases, then if there is no review inflation, Metacritic and Amazon scores should serve as comparable measures of quality, and the difference between the two should not vary in any systematic way. Thus, our dependent variable of interest will be defined as follows:

$$inflation_i \equiv metacritic_i - amazon_i \tag{8}$$

This variable serves as a proxy for η , the rate of review inflation. To test the two comparative statics described in the previous section, we require proxies for R (producer bargaining power) and θ (prior distribution on quality). For R , we will use firm *output*; specifically, we will use the number of unique games released by the producer in a given year. Both total advertising spending and potential content for reviewers increase with firm output, so it is reasonable to conclude that larger firms have greater bargaining power over reviewers. It follows that a positive relationship between *inflation* and *output* is consistent with our comparative static on R . Note specifically that we use number of games produced rather than sales revenue or sales volume to mitigate endogeneity concerns.

There is little evidence that θ varies by producer size. Like movie studios, even the largest, most successful game producers release a significant fraction of low-quality games. Indeed, the correlation between firm output and a game’s Amazon score is only about 0.01. Instead, consider splitting the review data into two groups: one group of games that are part of pre-existing game franchises and the other that are not. Define a dummy variable *sequel* that is equal to 1 for games in the former group. Sequels are not produced at random; there is a selection effect wherein only sufficiently successful games warrant the development of sequels. Since the quality of a sequel is positively correlated with the quality of its predecessor, we can conclude that expected quality of a random sequel is higher than that of a random non-sequel (indeed, sequels have an average Amazon rating about two points higher than non-sequels). It follows that a positive relationship between *inflation* and *sequel* is consistent with our comparative static on θ .

The model I estimate is a least squares regression with robust standard errors clustered by producer and year of release. Formally, the model takes the following form:

$$inflation_{ijt} = \delta_0 + \gamma_1 \times output_{jt} + \gamma_2 \times sequel_{ijt} + \mathbf{z}'_{ijt}\boldsymbol{\delta} + \epsilon_{ijt} \quad (9)$$

Where subscripts index observations by game i , producer j , and release year t . \mathbf{z} is a vector of control variables discussed in greater detail below. The baseline estimation is based on 550 observations of games released by publicly owned firms for the Xbox 360 console between 2007 and 2013.

Before analyzing the estimation results, it is necessary to discuss several important control variables. The largest concern is that user reviews on Amazon suffer from purchasing bias: only consumers who purchase the game write reviews, and only consumers who believe *ex*

ante that they will like the game will purchase it.¹⁸ On the other hand, expert reviewers are required to review whichever games come across their desks. The magnitude of this effect is likely larger for games of poorer quality; the worse a game’s reputation, the more likely only a person with an especially high interest would purchase it. To control for this effect, I include (the log of) the total number of Amazon ratings for each observation. Since the number of ratings correlates strongly with the game’s sales volume, a lower number of ratings corresponds with more severe purchasing bias.

I additionally include a dummy variable equal to 1 for games that are based on movies or television programs. Games based on other media tend to be panned by expert critics at a near-universal rate (the average Metacritic score for these games is nearly 15 points lower than the average for other games!). Average Amazon ratings between these two groups, on the other hand, are essentially identical. The concern regarding these games is twofold. First, purchasing bias in these cases may be especially pernicious. Because games based on other media have such a poor reputation for quality, the only consumers who buy them are those who enjoy them simply because of the nature of the content. Since only such people write user reviews on Amazon, they are a poor indicator of a game’s “objective” quality. Secondly, there may be other persistent taste differences. Games based on other media are disproportionately targeted at children. As a result, many of their reviews are written by parents or grandparents about how much the recipient of the game enjoys it. As with all media, there are surely taste differences between adult reviewers and child consumers.

Finally, it is of interest to examine whether firms put more pressure on reviewers to inflate reviews when they are struggling financially. To test this hypothesis, I include as an independent variable the percent change in producer stock price over the one-year window

¹⁸The existence of purchasing bias in user reviews is described and demonstrated elegantly by Hu et al. (2009). This article also examines the influence of reporting effects and concludes that, while reporting effects affect the variance of reviews, it is purchasing bias that is responsible for the asymmetry in the review distribution that biases the mean review score.

beginning three months prior to each game’s release (i.e. if a game was released on April 1, 2010, the relevant window is January 1, 2009 through January 1, 2010). For obvious reasons, I exclude observations for which there was a stock split during this window. These variables are summarized in Table 2.

Table 2: Summary Statistics for Variables

Variable	Mean	Std. Dev.	Min.	Max.
inflation	-7.931	13.227	-54.101	32.478
metacritic	70.269	14.899	22	98
amazon	78.200	9.927	41.429	100
output	10.136	6.414	1	24
sequel	.662	.474	0	1
#Amz_ratings	168.809	322.633	1	3006
movie_tv	.153	.360	0	1
%Δstock_price	.064	.473	-.896	4.272

The estimation results are displayed in Table 3. In the baseline regression, \mathbf{z} includes only the above three control variables. Table 3 also includes the results of several other specifications that are detailed in Appendix B.

Table 3: Estimation results

Variable	Baseline	Year effects	Geographic effects	No Microsoft	No congloms	Many reviews
output	.290*** (.089)	.203*** (.076)	.285*** (.105)	.290*** (.091)	.287*** (.087)	.371*** (.075)
sequel	3.143*** (.955)	3.411*** (.925)	3.200*** (.963)	3.044*** (.993)	3.276*** (1.060)	2.953*** (.995)
$\ln(\#Amz_ratings)$	3.874*** (.449)	4.607*** (.411)	3.839*** (.469)	3.815*** (.478)	3.971*** (.518)	3.589*** (.481)
movie_tv	-10.89*** (1.494)	-10.52*** (1.300)	-10.85*** (1.499)	-11.00*** (1.494)	-11.436*** (1.869)	-11.09*** (1.627)
%Δstock_price	-2.272* (1.179)	-.990 (1.295)	-2.151* (1.194)	-2.115* (1.200)	-1.995 (1.202)	-1.947** (.880)
Intercept	-27.51*** (2.338)	-33.75*** (3.435)	-27.30*** (2.871)	-27.14*** (2.386)	-27.854*** (2.573)	-26.17*** (2.486)

Significance levels: * = .10 | ** = .05 | *** = .01

We see that the number of Amazon reviews and the movie/TV dummy are extremely significant in the directions hypothesized. Specifically, more Amazon reviews indicate less purchasing bias and therefore higher Metacritic scores relative to Amazon scores. Additionally, a game based on a movie or television program was penalized approximately 11 points by experts relative to non-experts, though this difference is not distributed uniformly across genres. Even controlling for these significant effects, the relationships between the dependent variable and *output* and *sequel* are both positive and highly significant. This suggests (a) large, powerful producers are the beneficiaries of inflated reviews and (b) experts exhibit a conservatism bias with respect to games from pre-existing franchises. The finding of conservatism bias is consistent with the work of Camara and Dupuis (2014) on the film industry. Moreover, these results appear to be robust to a variety of specifications. For details on these robustness checks, see Appendix B.

The effect of recent stock price changes is inconclusive. In the baseline regression above, the effect is marginally significant and negative (suggesting firms whose stock price has fallen receive more heavily inflated reviews). However, the significance of this effect is not robust to changes in the window over which stock price change is considered. Depending on the window selected, significance varies from none ($p = .942$) to considerable ($p = .029$). See Appendix B for more details.

4 Equilibrium Inflation and Network Effects

To this point, we have seen that the mere existence of network effects can function as an on/off switch with respect to review inflation and that the video game journalism industry appears to be in an inflation equilibrium consistent with the predictions of that theory. The final piece of the puzzle is to understand how equilibrium inflation varies with the equilibrium size of the network externality. We can then exploit network effect heterogeneity in the video game market to rule out some alternative explanations for the existence of review inflation.

Recall from section 2.3 that reviewer revenue R and reviewer reputation cost ϕ both depend on early adoption premium λ_e .¹⁹ The natural relationship would be as follows: Larger network effects increase the demand for reviews but also the reputational damage that results from inaccurate reviews. We can therefore assume that $R'(\lambda_e) \geq 0$ and $\phi'(\lambda_e) \geq 0$. Since, in an inflation equilibrium, we can characterize equilibrium inflation by:

$$\eta^* = \frac{R(\lambda_e)}{(1 - \theta)\phi(\lambda_e)}$$

We can then show the following corollary:

Corollary 1 *The relationship between the equilibrium rate of inflation and the size of the early adoption premium depends on the relative elasticities of the reviewer’s revenue and reputation functions. Specifically:*

$$\frac{\partial \eta^*}{\partial \lambda_e} \begin{cases} \geq 0 & : \frac{R'(\lambda_e)}{R(\lambda_e)} \geq \frac{\phi'(\lambda_e)}{\phi(\lambda_e)} \\ < 0 & : \frac{R'(\lambda_e)}{R(\lambda_e)} < \frac{\phi'(\lambda_e)}{\phi(\lambda_e)} \end{cases}$$

The exact structure of the R and ϕ functions is beyond the scope of this paper but a potentially interesting topic for further study. For our purposes, we see that the relationship between network effects and inflation is theoretically ambiguous, and that a necessary and sufficient condition for this relationship to be positive is for the revenue effect to outweigh the reputational effect in equilibrium.

What does this mean for the video game industry? Quantifying the magnitude of network effects at the observation level would be difficult, but we can exploit heterogeneity at the genre level – some genres emphasize competition, collaboration, and multiplayer features more heavily than others. To test this theory, I add two additional dummy variables to the

¹⁹Depending on β instead would be equivalent for our purposes, although λ_e already depends on β in an inflation equilibrium since $\gamma_p = \beta$ when $m = H$.

regression specification from Section 3. The *shooter_sports* variable is equal to 1 for games from the shooter and sports game genres (approximately 36% of observations). The value of these games is weighted heavily on their multiplayer features. The *AA* variable is equal to 1 for games from the action-adventure genre (approximately 40% of observations). These games are disproportionately likely to be single-player. The results are displayed in Table 4.

Table 4: Estimation results : Genre Effects

Variable	Baseline	Year effects	Geographic effects	No Microsoft	No congloms	Many reviews
shooter_sports	3.842*** (1.069)	3.738*** (1.062)	3.878*** (1.074)	4.144*** (1.152)	4.355*** (1.162)	3.876*** (1.102)
AA	-3.361*** (1.111)	-2.982*** (1.122)	-3.396*** (1.108)	-3.206*** (1.163)	-2.777** (1.194)	-3.036** (1.224)
output	.198** (.083)	.124* (.069)	.224** (.097)	.190** (.085)	.195** (.083)	.276*** (.072)
sequel	2.150** (.938)	2.449*** (.910)	2.179** (.926)	2.055** (.990)	2.359** (1.074)	1.623* (.949)
$\ln(\#Amz_ratings)$	3.995*** (.432)	4.675*** (.403)	4.037*** (.438)	4.031*** (.457)	4.169*** (.501)	3.516*** (.444)
movie_tv	-7.768*** (1.518)	-7.668*** (1.381)	-7.618*** (1.511)	-7.764*** (1.515)	-8.209*** (1.899)	-8.600*** (1.658)
% Δ stock_price	-1.671 (1.156)	-.477 (1.239)	-1.665 (1.153)	-1.539 (1.170)	-1.480 (1.183)	-1.442* (.837)
Intercept	-26.95*** (2.366)	-32.69*** (3.670)	-27.68*** (2.777)	-27.12*** (2.428)	-28.17*** (2.603)	-24.44*** (2.375)

Significance levels: * = .10 | ** = .05 | *** = .01

We observe that inflation varies substantially between these two genre groups to the tune of over seven points on average. In the context of the current model, this suggests that the video game journalism industry is operating in an equilibrium where the marginal revenue effect exceeds the marginal reputational effect.

More importantly, however, this result rules out several alternative explanations for how inflation can arise in equilibrium. For example, naive consumer beliefs would give rise to inflation of the sort we observe in our baseline regression from Section 3, but consumer naivete fails to explain such large heterogeneity in inflation across genres.

5 Welfare Implications

5.1 How Competition Decreases Review Quality

Implicit in the earlier theoretical model is the assumption that the supply of reviewers is competitive. Conventional economic wisdom suggests that, if consumers value good information, then increased competition among reviewers should result in better information. However, if there are many different reviewers (as there are in the video game industry and many others), then the producer does not care much if it cannot contract with a specific reviewer and is therefore better positioned to leverage inflated reviews from each.

To see how this dynamic reduces review quality, first consider the counterfactual case of the monopolist reviewer and assume a more generalized bargaining process. As producer payoff strictly increases in η (in an inflation equilibrium) and reviewer payoff strictly decreases in η , we can assume that the bargaining result between the producer and the reviewer would be $\eta_1^* \in (0, \eta^*)$, where η^* is the value of η from our earlier equilibrium that binds the individual rationality constraint of the reviewer. The proximity of η_1^* to the endpoints of this range is determined by the parameters in this hypothetical model that characterize each firm's bargaining power. Consider, then, the case where there are two reviewers. Each individual reviewer has less bargaining power, and we would thus expect a bargaining outcome $\eta_2^* > \eta_1^*$. It is straightforward to see that, as the number of reviewers approaches perfect competition, the bargaining power of each individual reviewer converges to zero, and the equilibrium level of review inflation converges to η^* . Thus, increased competition amongst reviewers actually increases the pressure for each reviewer to cave to the demands of the producer and ultimately reduces the quality of information transmitted in equilibrium.

5.2 Stuck in a Bad Equilibrium

Empirical evidence from the previous section suggests that the video game industry is in an inflation equilibrium (type 3) of the sort described in Proposition 3. It is noteworthy, however, that this equilibrium is Pareto inferior to equilibrium type 1, where reviews are truthful and enthusiasts always early adopt. Producer profit is strictly higher in equilibrium 1 since the producer absorbs the surplus gain from fully exploiting the network externality.²⁰ Reviewer payoff is strictly higher because the reviewer suffers no reputational losses. Consumer expected payoff is zero in either equilibrium. This suggests that a policy intervention that pushes the video game industry from its current equilibrium into another could have a positive impact on welfare.

Germane to this idea is the increasing popularity of the “freemium” pricing structure in the software industry. Under the freemium structure, consumers are typically granted access to a software application for free. The producer of the software subsequently earns revenue from “microtransactions” – selling cosmetic and functional upgrades as well as additional content. In the context of our model, this should result in maximal exploitation of network effects since even non-enthusiasts are willing to early adopt at no cost.²¹ Some types of software, including more conventional stand-alone video games, are not as conducive to earning revenue from microtransactions; as a result, game publishers have not adopted the freemium format for many of their products. This suggests that subsidizing firms that adopt the freemium structure – a policy suggestion that might seem outlandish at face value – could indeed result in increased welfare (for political feasibility, such a subsidy could perhaps be funded by an excise tax on microtransactions).

²⁰Formally, producer profit in equilibrium 1 is $\pi^1 = \beta\lambda(\theta H + (1 - \theta)L) + (1 - \beta)(\theta H + (1 - \theta)L)$. This number is strictly greater than $\pi^3 = \beta\lambda(\theta H + (1 - \theta)\eta L) + \beta(1 - \theta)(1 - \eta)L + (1 - \beta)(\theta H + (1 - \theta)L)$.

²¹Technically, there remains the opportunity cost of time spent trying out the product. However, casual observation of the market for apps and the adoption rate differential between free apps and costly apps suggests that consumers are far more willing to invest time to try a product than money. Moreover, since product reviews are partially informative even in an inflation equilibrium, they will still have value in reducing search frictions with respect to these opportunity costs.

6 Conclusion

The analysis in this paper yields several new insights about how expert product reviews affect the market for their products. First, in line with conventional economic wisdom, it is not obvious that a firm can use inflated reviews to increase its profits when the review acts solely as a mechanism for information transmission. However, if a positive review actually creates economic value in some way – even if both the size of that value and the fraction of people who consume it are arbitrarily small – then it is optimal for a producer to inflate reviews to the greatest extent possible. The extent to which review inflation is possible unsurprisingly depends on how much bargaining power a producer has over its reviewers; the punishment for not cooperating with the producer must exceed the reputational losses that result from inflated reviews.

One theory of how positive reviews can create value follows from the assumption that at least some portion of consumers benefits from positive network effects. Since virtually any form of entertainment or hobby can generate a positive network effect between consumers who have consumed the same media or participated in the same activity via social interaction, this assumption is quite reasonable. The result is a coordination subgame among such consumers; to maximize the benefit of the externality, they prefer to make purchase decisions in unison. Reviews can serve as a coordination mechanism for consumers who might not otherwise be able to communicate. In particular, expert reviews that appear before the release of the product (and before non-experts are able to disseminate their opinions on the internet) are of value to those who must decide whether to early adopt. This theory provides a possible explanation to an empirical puzzle in the marketing literature: although the prevalence of user reviews and digital word-of-mouth has sharply increased in the digital age, the impact of expert reviews on sales persists. If expert reviews served no function but to transmit information, this would not be the case.

The theory in this article provides several comparative statics that are testable using review data. Such data from the video game industry show evidence of review inflation. This finding is itself of interest, as those close to the industry have long suspected the existence of bias in video game reviews. Also of interest, however, is the way in which review inflation manifests; specifically, the extent of inflation increases with the size of the producer of the game, and expert reviewers appear to exhibit conservatism bias for games that are new entries in pre-existing franchises. Both of these findings are consistent with the predictions of the theoretical model. Additionally, inflation is higher for games in genres that feature strong network effects. This result is consistent with the network effects theory of inflation but inconsistent with some alternative theories.

These results have some intriguing welfare implications for the video game industry. First, despite the fact that both reviewers and consumers prefer truthful information over less truthful information, we find that increased competition in the market for reviews actually decreases the quality of information transmitted. This suggests that a single monopolist reviewer could maximize the quality of information transmitted, though this depends on the assumption that such a reviewer is motivated solely by profit. Second, the empirical analysis suggests that the video game industry is stuck in a Pareto-dominated equilibrium. This is because the positive network externalities in this case are a free lunch that is not being fully exploited. It follows that policy aimed at maximally exploiting the positive network effects in this industry could increase welfare.

The next logical step would be empirical analysis of other industries to see how consistently review inflation varies with the predictions of the model. In the film industry, for example, we would expect to find conservatism bias as Camara and Dupuis (2014) do, but we would expect the variation based on studio size to be weak because movie studios have little bargaining power over the general interest publications that publish reviews of their products. As research in this area is mostly a post-internet phenomenon, the literature is fairly small, and

there is much more work yet to be done.

Appendix A: Proofs

Proof of Proposition 1

Proof. It is established in the main text that, in any equilibrium, enthusiasts will early adopt at price $p^* = \mathbb{E}[q|m]$ and non-enthusiasts will wait and purchase at price $w^* = q$. What remains is to solve for the producer's profit-maximizing editorial policy subject to the individual rationality constraint of the reviewer. Given p^* , w^* , and some editorial policy (η, ν) , ex ante expected firm profit is:

$$\pi = \beta [Pr(m = H)\lambda_e\mathbb{E}[q|m = H] + Pr(m = L)\lambda_e\mathbb{E}[q|m = L]] + (1 - \beta)\mathbb{E}[q] \quad (10)$$

Since:

$$Pr(m = H) = \theta(1 - \nu) + (1 - \theta)\eta \quad (11)$$

$$Pr(m = L) = \theta\nu + (1 - \theta)(1 - \eta) \quad (12)$$

$$\mathbb{E}[q|m = H] = \frac{\theta(1 - \nu)H + (1 - \theta)\eta L}{\theta(1 - \nu) + (1 - \theta)\eta} \quad (13)$$

$$\mathbb{E}[q|m = L] = \frac{\theta\nu H + (1 - \theta)(1 - \eta)L}{\theta\nu + (1 - \theta)(1 - \eta)} \quad (14)$$

$$\mathbb{E}[q] = \theta H + (1 - \theta)L \quad (15)$$

Expected profit reduces to:

$$\pi = \beta\lambda_e[\theta(1 - \nu)H + (1 - \theta)\eta L + \theta\nu H + (1 - \theta)(1 - \eta)L] + (1 - \beta)(\theta H + (1 - \theta)L) \quad (16)$$

$$= \beta\lambda_e[\theta H + (1 - \theta)L] + (1 - \beta)(\theta H + (1 - \theta)L) \quad (17)$$

The above expression is independent of (η, ν) and also exactly equal to firm profit if there is no review (in which case the firm sells to enthusiasts at early adoption price $p = \lambda_e \mathbb{E}[q]$ and to non-enthusiasts at wait price $w = q$). Since producer payoff is independent of η and ν , any (η, ν) that satisfies individual rationality for the reviewer can be an equilibrium. ■

Proof of Proposition 2

Proof. It is established in the main text that, in any equilibrium, when $m = L$, the best the producer can charge to either consumer is $p = \mathbb{E}[q|m = L]$ in the early adoption stage and $w = q$ in the wait stage. Since the producer is risk neutral, these are effectively equivalent. When $m = H$, enthusiasts early adopt at price $\lambda_e(H)\mathbb{E}[q|m = H]$ and non-enthusiasts purchase at wait price $w = q$. We can now characterize producer profit for a given (η, ν) :

$$\pi = Pr(m = H)(\beta\lambda_e(H)\mathbb{E}[q|m = H]) + Pr(m = L)(\beta\mathbb{E}[q|m = L]) + (1 - \beta)\mathbb{E}[q] \quad (18)$$

Since, as in the proof of Proposition 1, we have:

$$Pr(m = H) = \theta(1 - \nu) + (1 - \theta)\eta \quad (19)$$

$$Pr(m = L) = \theta\nu + (1 - \theta)(1 - \eta) \quad (20)$$

$$\mathbb{E}[q|m = H] = \frac{\theta(1 - \nu)H + (1 - \theta)\eta L}{\theta(1 - \nu) + (1 - \theta)\eta} \quad (21)$$

$$\mathbb{E}[q|m = L] = \frac{\theta\nu H + (1 - \theta)(1 - \eta)L}{\theta\nu + (1 - \theta)(1 - \eta)} \quad (22)$$

$$\mathbb{E}[q] = \theta H + (1 - \theta)L \quad (23)$$

Expected profit reduces to:

$$\pi = \beta\lambda_e(H)(\theta(1 - \nu)H + (1 - \theta)\eta L) + \beta(\theta\nu H + (1 - \theta)(1 - \eta)L) + (1 - \beta)(\theta H + (1 - \theta)L) \quad (24)$$

$$= \beta\theta H(\lambda_e(H) + (1 - \lambda_e(H))\nu) + \beta(1 - \theta)L(1 + (\lambda_e(H) - 1)\eta) + (1 - \beta)(\theta H + (1 - \theta)L) \quad (25)$$

Since $\lambda_e(H) > 1$, it is immediate that producer profit is strictly increasing in η and strictly decreasing in ν . Thus, the producer always proposes $\nu^* = 0$ and chooses the highest possible value for η , i.e. the value that binds the individual rationality constraint of the reviewer. Thus, η^* is characterized by:

$$R - [(1 - \theta)\eta^* + \theta\nu^*]\phi = 0 \quad (26)$$

$$\iff \eta^* = \frac{R}{(1 - \theta)\phi} > 0 \quad (27)$$

■

Proof of Proposition 3

Proof. As always, we have $\bar{w} = q$ for any type of consumer. The remainder of the equilibrium, including the producer's optimal editorial policy and early adoption price, will depend on how enthusiasts coordinate in the early adoption stage.

The proof proceeds in two steps. First, we will show that, in equilibrium, every enthusiast must take the same action. This will result in five possible coordinated strategies for enthusiasts. In the second step, we construct the essentially unique equilibrium for each of the five cases.

Step 1: Every enthusiast chooses the same action. By contradiction, suppose that some mass $\alpha \in (0, \beta)$ of enthusiasts chooses to early adopt in equilibrium. If this is the result of an optimal strategy, it must also be that $\bar{p} = \lambda_e(\alpha)\mathbb{E}[q|m]$. Then, consider an enthusiast who chose not to early adopt. The payoff for this consumer is zero. But because of that, this consumer is strictly better off making a purchase at any $\hat{p} < \lambda_e(\alpha)\mathbb{E}[q|m]$, which implies $\bar{p} \geq \lambda_e(\alpha)\mathbb{E}[q|m]$ for this consumer. It is thus optimal for this consumer to early adopt, a contradiction. It follows that, in any equilibrium outcome, either every enthusiast early adopts or none does.

Step 2: Constructing equilibria. It follows from Step 1 that enthusiasts can coordinate in equilibrium in five ways:²²

1. Enthusiasts early adopt regardless of the review.
2. Enthusiasts do not early adopt regardless of the review.
3. Enthusiasts early adopt if and only if $m = H$.

²²It is additionally possible that enthusiasts early adopt if and only if $m \in \{H, \emptyset\}$ or $m \in \{L, \emptyset\}$. However, as described in the proof, $m = \emptyset$ is off-path for these equilibria, making them equivalent to types 3 and 4, respectively.

4. Enthusiasts early adopt if and only if $m = L$.

5. Enthusiasts early adopt if and only if $m = \emptyset$.

Enthusiasts always early adopt. Since β enthusiasts early adopt, it follows that $\bar{p} = \lambda_e(\beta)\mathbb{E}[q|m]$ and that $p^* = \lambda_e(\beta)\mathbb{E}[q|m]$. To find the optimal editorial policy, the producer maximizes:

$$\pi = \beta\lambda_e(\beta)[\theta(1 - \nu)H + (1 - \theta)\eta L + \theta\nu H + (1 - \theta)(1 - \eta)L] + (1 - \beta)(\theta H + (1 - \theta)L) \quad (28)$$

This is exactly the expression from the proof of Proposition 1. Thus, producer profit is independent of its editorial policy.

Enthusiasts never early adopt. For this case to be an equilibrium, it must be that $p^* > \bar{p} = \mathbb{E}[q|m]$. If $\bar{p} > \mathbb{E}[q|m]$, then it is optimal for the producer to choose $p = \bar{p}$ and enthusiasts subsequently early adopt. Since $\bar{p} < \mathbb{E}[q|m]$ is never optimal for consumers, the equality follows. Since the producer is indifferent between selling at $p = \mathbb{E}[q|m]$ and $w = q$, it follows that $p > \mathbb{E}[q|m]$ can be optimal in the early adoption stage. Finally, since all uncertainty is resolved prior to any purchase decisions, the editorial policy is obviously irrelevant.

Enthusiasts early adopt if and only if $m = H$. In this case, \bar{p} depends on m . Specifically, we have $\bar{p}(H) = \lambda_e(\beta)\mathbb{E}[q|H]$ and $\bar{p}(L) = \mathbb{E}[q|L]$. The producer therefore maximizes:

$$\pi = Pr(m = H)(\beta\lambda_e(\beta)\mathbb{E}[q|m = H]) + Pr(m = L)(\beta\mathbb{E}[q|m = L]) + (1 - \beta)\mathbb{E}[q] \quad (29)$$

$$= \beta\lambda_e(\beta)(\theta(1 - \nu)H + (1 - \theta)\eta L) + \beta(\theta\nu H + (1 - \theta)(1 - \eta)L) + (1 - \beta)(\theta H + (1 - \theta)L) \quad (30)$$

This is exactly the expression from the proof of Proposition 2. Thus, the optimal editorial policy is maximal inflation.

Enthusiasts early adopt if and only if $m = L$. In this case, \bar{p} again depends on m . Specifically, we have $\bar{p}(L) = \lambda_e(\beta)\mathbb{E}[q|L]$ and $\bar{p}(H) = \mathbb{E}[q|H]$. The producer therefore maximizes:

$$\pi = Pr(m = H)(\beta\mathbb{E}[q|m = H]) + Pr(m = L)(\beta\lambda_e(\beta)\mathbb{E}[q|m = L]) + (1 - \beta)\mathbb{E}[q] \quad (31)$$

And, as in prior propositions, we have:

$$Pr(m = H) = \theta(1 - \nu) + (1 - \theta)\eta \quad (32)$$

$$Pr(m = L) = \theta\nu + (1 - \theta)(1 - \eta) \quad (33)$$

$$\mathbb{E}[q|m = H] = \frac{\theta(1 - \nu)H + (1 - \theta)\eta L}{\theta(1 - \nu) + (1 - \theta)\eta} \quad (34)$$

$$\mathbb{E}[q|m = L] = \frac{\theta\nu H + (1 - \theta)(1 - \eta)L}{\theta\nu + (1 - \theta)(1 - \eta)} \quad (35)$$

$$\mathbb{E}[q] = \theta H + (1 - \theta)L \quad (36)$$

Expected profit reduces to:

$$\pi = \beta(\theta(1 - \nu)H + (1 - \theta)\eta L) + \beta\lambda_e(\beta)(\theta\nu H + (1 - \theta)(1 - \eta)L) + (1 - \beta)(\theta H + (1 - \theta)L) \quad (37)$$

$$= \beta\theta H(1 + (\lambda_e(\beta) - 1)\nu) + \beta(1 - \theta)L(\lambda_e(\beta) + (1 - \lambda_e(\beta))\eta) + (1 - \beta)(\theta H + (1 - \theta)L) \quad (38)$$

Since $\lambda_e(\beta) > 1$, it is immediate that producer profit is strictly increasing in ν and strictly decreasing in η . Thus, the producer always proposes $\eta^* = 0$ and chooses the highest possible value for ν , i.e. the value that binds the individual rationality constraint of the reviewer. Thus, ν^* is characterized by:

$$R - [(1 - \theta)\eta^* + \theta\nu^*]\phi = 0 \quad (39)$$

$$\iff \nu^* = \frac{R}{\theta\phi} > 0 \quad (40)$$

Enthusiasts early adopt if $m = \emptyset$. If enthusiasts early adopt when $m = \emptyset$, it follows that $p^* = \bar{p} = \lambda_e(\beta)\mathbb{E}[q] = \theta H + (1 - \theta)L$. There are three possible cases here not covered by the previous equilibria.

If enthusiasts early adopt if and only if $m = \emptyset$, it is immediate that the optimal editorial policy is $m = \emptyset$ since $[\beta\lambda_e(\beta) + (1 - \beta)](\theta H + (1 - \theta)L) > \theta H + (1 - \theta)L$. This is the fifth type of equilibrium that is given in the proposition.

If enthusiasts early adopt if and only if $m \in \{H, \emptyset\}$, then the optimal editorial policy is the same as in Equilibrium 3. To see that this is true, notice that producer profit when $m = \emptyset$ is exactly equal to the case of Equilibrium 3 when $\eta = \nu = 0$. But we already know that the producer can do strictly better by choosing $\eta > 0$, so $m = \emptyset$ will always be off-path

in this case.

If enthusiasts early adopt if and only if $m \in \{L, \emptyset\}$, then the optimal editorial policy is the same as in Equilibrium 4. To see that this is true, notice that producer profit when $m = \emptyset$ is exactly equal to the case of Equilibrium 4 when $\eta = \nu = 0$. But we already know that the producer can do strictly better by choosing $\nu > 0$, so $m = \emptyset$ will always be off-path in this case. ■

Proof of Corollary 1

Proof. Since the equilibrium rate of inflation is given by:

$$\eta^* = \frac{R(\lambda_e)}{(1 - \theta)\phi(\lambda_e)} \quad (41)$$

It follows that:

$$\frac{\partial \eta^*}{\partial \lambda_e} = \frac{(1 - \theta)\phi(\lambda_e)R'(\lambda_e) - R(\lambda_e)(1 - \theta)\phi'(\lambda_e)}{[(1 - \theta)\phi(\lambda_e)]^2} \quad (42)$$

This expression is weakly positive if and only if:

$$\phi(\lambda_e)R'(\lambda_e) - R(\lambda_e)\phi'(\lambda_e) \geq 0 \quad (43)$$

$$\iff \frac{R'(\lambda_e)}{R(\lambda_e)} \geq \frac{\phi'(\lambda_e)}{\phi(\lambda_e)} \quad (44)$$

Where the final inequality holds because $R(\lambda_e)$ and $\phi(\lambda_e)$ are strictly positive. ■

Appendix B: Data Details and Robustness Checks

The data used in this analysis are based on 550 games released for the Xbox 360 video game console between 2007 and 2013. Summary statistics for the variables used in the main regression are presented in Table 2. I subsequently discuss each variable independently.

inflation. The *inflation* variable is the simple difference between the *metacritic* and *amazon* variables described below. As we can see from Table 2, the word “inflation” is a bit of a misnomer; upward purchasing bias in user reviews results in most *inflation* values being negative. Of course, we care only about the relative difference between observations, so it does not matter if the values for this variable are negative.

metacritic. Metacritic, founded in 2001, is a website that aggregates expert reviews using different scales of measurement into a single numerical score between 0 and 100. The precise aggregation algorithm is proprietary; however, Metacritic score data has been increasingly used in empirical work on expert reviews and has become accepted as a valid measure of expert sentiment. The scores in the sample range from 22 to 98, so there is no concern about censoring. That the lowest score is 22 is helpful; when normalizing Amazon ratings (below) to a comparable scale, the lowest possible rating (one star) is transformed into a score of 20. Thus, both ratings systems effectively use a scale of 20-100. The number of individual reviews on which scores in the sample were based range from 3 to 100. As I demonstrate below, dropping observations based on having “too few” individual reviews does not impact the results. The scores used in the current analysis were gathered manually in January 2015.

amazon. This variable consists of each observation’s average Amazon rating (out of five stars) multiplied by 20 to get a number between 20 and 100. That the mean is higher than that of the *metacritic* variable is attributed to purchasing bias, as discussed elsewhere in the article. The lowest score was 41.4, and there were only two observations with perfect scores of 100, so censoring is again of minimal concern. The number of individual ratings (which is

itself an independent variable in the model discussed below) ranged from 1 to 3006 (median 60). As with *metacritic*, dropping observations with insufficiently many individual ratings does not affect the results. The ratings used in the current analysis were gathered manually in January 2015 and may therefore vary slightly from current ratings.

output. This variable measures how many individual games were released by a particular producer in a given year. Specifically, it is the number of other games also included in this sample, i.e. it does not include games produced for other consoles, mobile devices, or personal computers. Every producer in the sample (except Microsoft, discussed further below) produces games for multiple platforms, so the number of games produced for Xbox 360 should correlate strongly with the number of total games produced, which is intended as a proxy for bargaining power over reviewers. The values for this variable range from 1 to 24.

sequel. This variable is a dummy equal to 1 if the observation is part of a pre-existing game franchise. Nearly two-thirds of observations qualify as sequels according to this standard. This finding is unsurprising, as sequels are low hanging fruit for producers; they cost less to produce than a game developed from scratch, and they are less risky because they already have an established fanbase.

#Amz_ratings. This variable is the total number of individual ratings for the observation on Amazon at the time of retrieval in January 2015. As mentioned above and discussed further below, values for this variable range from 1 to 3006 (median 60), and restricting the regression to observations with a minimum number of ratings does not affect results.

movie_tv. This variable is a dummy equal to 1 if the game is based directly on a pre-existing movie or television program. Its value is equal to 1 for approximately 15 percent of observations.

% Δ stock_price. This variable is the percent change in producer stock price from fifteen months prior to a game’s release to three months prior. As discussed below, the choice of window here makes a substantial difference in the significance of this variable’s coefficient but not in the significance of other variables in the regression. Data were retrieved from Yahoo! Finance in May 2015. Stock prices used are the “adjusted closing price” that accounts for dividend payments. The regression omits any observations for which this window includes a stock split, which reduced the number of observations from 640 to 550.

The Effects of Stock Price Changes Over Different Windows

The idea behind including stock price changes in the regression is the hypothesis that producers who are struggling financially might put increased pressure on reviewers to inflate reviews. This is an especially relevant notion because the range of dates from which data is used (2007 to 2013) includes a major recession and subsequent recovery. A decomposition of results by year (see the next subsection) is consistent with the idea that the recession led firms to increase pressure on reviewers; however, analysis with actual stock data is less conclusive.

The natural starting point for incorporating these effects is three months prior to a game’s release. At this point, promotional campaigns are shifting into overdrive, but hype created by expert reviews themselves (which are often released just days or weeks prior to the game itself) have not yet had an opportunity to affect stock price. Thus, it is natural to include the producer’s financial performance over the previous year as a possible determinant of review inflation.

In the main regression, the coefficient on this variable is marginally significant. I show in this subsection, however, that this result can be decomposed into effects from two mutually exclusive partitions: stock price changes between three and nine months prior to release have virtually no effect (Table 5; $p = .739$), while stock price changes between nine and

fifteen months prior to release have a substantial effect in the expected direction, i.e. poor performance correlates with more inflation (Table 6; $p = .044$).

Table 5: Estimation results : 3-9 months

Variable	Coefficient	(Std. Err.)
output	0.295***	(0.095)
sequel	3.120***	(0.957)
$\ln(\#Amz_ratings)$	3.877***	(0.461)
movie_tv	-10.846***	(1.499)
$\% \Delta stock_price[3-9]$	-0.492	(1.473)
Intercept	-27.688***	(2.372)

Significance levels: * = .10 | ** = .05 | *** = .01

Table 6: Estimation results : 9-15 months

Variable	Coefficient	(Std. Err.)
output	0.290***	(0.090)
sequel	3.105***	(0.955)
$\ln(\#Amz_ratings)$	3.873***	(0.454)
movie_tv	-10.963***	(1.493)
$\% \Delta stock_price[9-15]$	-3.295**	(1.615)
Intercept	-27.482***	(2.348)

Significance levels: * = .10 | ** = .05 | *** = .01

It is not clear what is driving this result. One conjecture is that, nine months prior to release, most games have alpha builds that can be used for demonstrations, and this is when the promotional cycle and interaction with reviewers begins. However, promotional campaigns begin that early in earnest for only a small fraction of highly anticipated games. This is an area where further research may be of interest.

Year Effects

It is of interest to observe how the results of the main regression vary over time, particularly in the context of the major financial crisis of 2007-2009. Table 7 shows estimation results of the same regression with year dummies added (2013 is the base year).

Table 7: Estimation results : Year Effects

Variable	Coefficient	(Std. Err.)
output	0.203***	(0.076)
sequel	3.411***	(0.925)
$\ln(\#\text{Amz_ratings})$	4.607***	(0.411)
movie_tv	-10.518***	(1.300)
$\%\Delta\text{stock_price}$	-0.990	(1.295)
2007	6.589**	(3.062)
2008	5.262*	(3.079)
2009	8.954**	(3.599)
2010	1.908	(2.928)
2011	3.352	(2.757)
2012	-0.362	(3.200)
Intercept	-33.752***	(3.435)

Significance levels: * = .10 | ** = .05 | *** = .01

It is clear that review inflation was higher in the earlier years of the sample, apparently peaking in 2009. This is consistent with the theory that the recession caused producers to increase pressure on reviewers to inflate reviews. Tests of this theory with stock prices, however, yield results that are inconsistent at best.

An alternative explanation is that, in earlier years, the success of the Xbox 360 console remained uncertain. Thus, there was more at stake in pressuring reviewers to inflate reviews. By the latter half of the sample, the Xbox 360 was an established success; its successor, the Xbox One, was released in November 2013, and it is plausible that attention shifted away from the 360 as it simply became a less relevant platform.

Geographic Effects

In an earlier section, I detailed an allegation of journalistic corruption between a Japanese game producer and a Japanese video game magazine. It is worthwhile to investigate how the results of the main regression vary by country of origin. There are several reasons why this might matter. First, a Japanese producer might focus its corruption efforts, if any, on Japanese review publications that are underrepresented in the Metacritic aggregation

mechanism. Moreover, a fraction of Japanese games are never released in North America at all, so efforts to influence reviewers who target North American readers would be less valuable. In addition, cultural differences with respect to business practices could affect producer actions. To test if country (or continent) of origin is driving some or all of our results, I estimate the original regression with dummies for games that were produced by Japanese (25 percent of observations) or European (11.7 percent) producers. The results, displayed in Table 8, show no evidence that geographic origin is relevant.

Table 8: Estimation results : Geographic Effects

Variable	Coefficient	(Std. Err.)
output	0.285***	(0.105)
sequel	3.200***	(0.963)
$\ln(\#Amz_ratings)$	3.839***	(0.469)
movie_tv	-10.845***	(1.499)
$\% \Delta stock_price$	-2.151*	(1.194)
japan	-0.827	(1.528)
europe	1.269	(2.300)
Intercept	-27.295***	(2.871)

Significance levels: * = .10 | ** = .05 | *** = .01

Testing for Heterogeneous Target Audiences

One possible concern with the empirical analysis presented in this article is that different expert reviewers might have different target audiences, and this heterogeneity might drive some of the variation in the *inflation* variable. To test this hypothesis, I collected and compared review scores from two specific video game websites.

The first website is IGN, the most popular video game website on the internet as measured by monthly traffic.²³ Perhaps more importantly, IGN (or chief rival GameSpot) is typically the first result to pop up after a Google search for “review of game X.” Thus, of all expert reviewers, we can expect that IGN is the most heavily slanted toward casual gamers.

²³As of June 2015, IGN claims 68 million unique users and 513 million page views per month.

The second website is GamesRadar, a medium-sized site still large enough to provide reviews for nearly all released games. Precise figures on internet traffic vary, but estimates indicate that GamesRadar gets at most about 10 percent of the traffic that IGN does, and perhaps even less than that. More importantly, GamesRadar reviews do not appear near the top of Google results when searching for game reviews. Thus, at least on the margin, we expect that GamesRadar is more heavily slanted toward enthusiasts than IGN is.

Since both sites rate games on a scale from 1 to 10, their scores are directly comparable. Using review scores on 454 Xbox 360 games for which both sites have reviews, I estimate the following model:

$$IGN_{ijt} = \delta GR_{ijt} + \epsilon_{ijt} \tag{45}$$

If there are no systematic differences between the two sites' review generation processes, then we expect δ to be equal to 1 and idiosyncratic taste differences to be captured by the white noise term. Indeed, the estimate for δ is .99987 and we fail to reject the hypothesis that it is equal to 1 ($p = .9829$). I thus find no reason to believe there are systematic differences between reviewers based on target audience or any other criteria by which IGN and GamesRadar differ.

Dropping Microsoft Observations

Of 550 observations in the main regression, 41 are games produced by Microsoft. Microsoft is unique among producers in the sample in two ways. First, it is the only one that produces hardware in addition to software; thus, it could have more bargaining power than its *output* variable suggests. Second, it is the only producer for which all of its console games are Xbox-exclusive. While for all other producers, number of Xbox games produced is a reasonable proxy for total number of console games produced, this is not the case for Microsoft. For

this reason, it could actually have less bargaining power than its *output* variable suggests. To ensure that the presence of this unique firm is not altering the results, I estimate the regression again using only the 509 non-Microsoft observations. The results, displayed in Table 9, indicate that this exclusion makes little material difference.

Table 9: Estimation results : No Microsoft

Variable	Coefficient	(Std. Err.)
output	0.290***	(0.091)
sequel	3.044***	(0.993)
$\ln(\#\text{Amz_ratings})$	3.815***	(0.478)
movie_tv	-11.000***	(1.494)
$\%\Delta\text{stock_price}$	-2.115*	(1.200)
Intercept	-27.142***	(2.386)

Significance levels: * = .10 | ** = .05 | *** = .01

Dropping Observations from Media Conglomerates

The main regression makes use of stock price changes as a measure of financial performance. Three of the producers in the sample (Disney Interactive Studios, Microsoft Game Studios, and Warner Bros Interactive) are relatively small segments of huge media conglomerates (Walt Disney Company, Microsoft, and Time Warner, respectively). Stock price of the parent company may therefore not be reflective of the segment that produces video games. To test if this is affecting our results, I estimate the main regression model without the games produced by these three firms (470 remaining observations). The results, displayed in Table 10, indicate that this exclusion makes little material difference.

Dropping Observations with Few Reviews

Finally, there may be concern that observations with very few reviews (either Metacritic or Amazon) might have a disproportionate effect on results even though the number of Amazon reviews is already included as an independent variable to control for purchasing bias. Observations with a small number of reviews are not only vulnerable to manipulation

Table 10: Estimation results : No Conglomerates

Variable	Coefficient	(Std. Err.)
output	0.287***	(0.087)
sequel	3.276***	(1.060)
$\ln(\#Amz_ratings)$	3.971***	(0.518)
movie_tv	-11.436***	(1.869)
% Δ stock_price	-1.995	(1.202)
Intercept	-27.854***	(2.573)

Significance levels: * = .10 | ** = .05 | *** = .01

but also may simply not be representative of consumer sentiment regarding a product. To test whether this is the case, I estimate the main regression while restricting to observations with at least 15 Metacritic reviews (Table 11; 477 observations), at least 20 Amazon reviews (Table 12; 446 observations), or both (Table 13; 414 observations).

We see that these restrictions actually strengthen the confidence of the relationship between *inflation* and *output* while leaving the others relatively unchanged. I thus conclude that this is not a reason for concern.

Table 11: Estimation results : #MC_reviews \geq 15

Variable	Coefficient	(Std. Err.)
output	0.281***	(0.076)
sequel	3.029***	(0.957)
$\ln(\#Amz_ratings)$	3.062***	(0.392)
movie_tv	-10.791***	(1.377)
% Δ stock_price	-1.747**	(0.871)
Intercept	-22.689***	(2.216)

Significance levels: * = .10 | ** = .05 | *** = .01

Table 12: Estimation results : $\#Amz_ratings \geq 20$

Variable	Coefficient	(Std. Err.)
output	0.401***	(0.084)
sequel	3.214***	(1.140)
$\ln(\#Amz_ratings)$	3.914***	(0.451)
movie_tv	-11.583***	(1.614)
$\% \Delta stock_price$	-1.894*	(0.980)
Intercept	-28.714***	(2.456)

Significance levels: * = .10 | ** = .05 | *** = .01

Table 13: Estimation results : Both

Variable	Coefficient	(Std. Err.)
output	0.371***	(0.075)
sequel	2.953***	(0.995)
$\ln(\#Amz_ratings)$	3.589***	(0.481)
movie_tv	-11.088***	(1.627)
$\% \Delta stock_price$	-1.947**	(0.880)
Intercept	-26.166***	(2.486)

Significance levels: * = .10 | ** = .05 | *** = .01

Appendix C: Theory Robustness

The theoretical examples in the main body of this article are presented in an overly stylized way by design; it is intended to maximize the intuition of the reader. In this appendix, I explore the consequences of relaxing several assumptions.

The Case of $N \geq 2$ Reviewers

In reality, a producer interacts with many reviewers that review its products. This is reflected in the model by the fact that the supply of reviewers is implicitly treated as competitive. What happens if we assume $N \geq 2$ reviewers, each of which interacts simultaneously and independently with the producer? If we allow the producer to correlate the randomization mechanism so that all reviewers inflate the same products, then the entire model and all results remain exactly unchanged. Let us instead consider the case where the randomization mechanism is uncorrelated.

Increasing N improves the amount of information transmitted in an inflation equilibrium.²⁴ This result is due to the fact that $\nu^* = 0$ in an inflation equilibrium. Thus, if consumers observe even a single $m = L$, they know that $q = L$ as well. However, if $q = H$, then we will observe $m = H$ for all N reviewers, and consumer confidence in this signal will increase. As N trends to infinity, information transmission trends to perfect information. Expected profit for the producer decreases with N .²⁵

However, as long as N is finite, analogs of the three propositions from the main body of the article still hold. In Example 1, the extension is immediate; the combined signal of all reviewers simply replaces the single signal, and everything remains the same. In Example 2, if $\lambda_e(H)$ weakly increases in the number of H signals (so long as it strictly increases in at least one case), then the producer still has incentive to maximize inflation. In Example 3, the number of possible adoption rules increases exponentially as N does, but any equilibrium where early adoption depends on positive reviews (e.g. enthusiasts early adopt if there are at least $n < N$ positive reviews) will result in our familiar inflation equilibrium and the comparative statics that follow from it.

Producer Knowledge of Product Quality

In the model, the producer has no private knowledge of the quality of its product. If the producer has private information about quality, this may affect its optimal pricing policy, which in turn may allow consumers to infer information about quality.

To see how this is true, suppose that the producer knows the value of q with certainty. In our equilibria, the producer chooses $p^* = \lambda_e E[q|m]$. If this value is weakly greater than H , the private information does not matter. If $p^* < H$, however, then a producer who

²⁴This is not in conflict with idea that quality of information may decrease with N from section 4. There, monopolist and oligopolist reviewers had non-zero bargaining power that decreases in N . Here, reviewers have no bargaining power throughout.

²⁵This is immediate from the fact that the producer could have chosen truthful revelation in the first place but was strictly better off with maximal inflation.

knows $q = H$ would never sell at this price; it would strictly prefer to sell at $w = H$ in the wait stage. Since the producer has no means of credibly transmitting information beyond m , however, consumers will not be willing to early adopt at $p = H$. Since it cannot be in equilibrium that consumers early adopt at different prices depending on q but holding m constant,²⁶ it follows that consumers never early adopt at a price higher than $\mathbb{E}[q|m]$.

The existence of our inflation equilibrium thus requires that $\lambda_e \mathbb{E}[q|m] \geq H$.²⁷ Expansion of this inequality yields:

$$\lambda_e \cdot \frac{\theta(1 - \nu^*)H + (1 - \theta)\eta^*L}{\theta(1 - \nu^*) + (1 - \theta)\eta^*} \geq H \quad (46)$$

$$\iff \lambda_e \geq \frac{\theta\phi H + RH}{\theta\phi H + RL} > 1 \quad (47)$$

Thus, allowing the producer private knowledge of product quality weakens our result; whereas before λ_e could be arbitrarily close to 1, this is no longer the case. However, at least in the video game industry, it is very likely that λ_e indeed satisfies this condition. The equilibrium that follows when this condition is not satisfied involves no early adoption (obviously not true for video games) and no review inflation (inconsistent with the empirical findings of this article). Moreover, reviews themselves would be entirely superfluous (inconsistent with empirical work that shows expert reviews affect sales).

The Case of $N > 2$ Consumer Types

In the model, it is assumed that there are only two types of consumers, and that one of these types has the same payoff function regardless of whether or not he early adopts. This

²⁶If $p_H > p_L$, then the producer has incentive to charge p_H even when $q = L$, since consumers cannot discern the difference. If consumers are willing to pay $p_L > p_H$ even when they can tell from the price that $q = L$, then the producer is strictly better off increasing p_H to at least p_L .

²⁷It is perhaps most realistic that the producer observes a private noisy signal about the quality of its product rather than having full knowledge. If the producer observes signal s , then the same analysis applies with $\mathbb{E}[q|m, s]$ on the right hand side rather than H .

assumption drastically simplifies the optimal pricing problem for the producer, but it is not necessary for the important results in this article.

To see why this is true, suppose instead that there is a continuum of consumer types with distribution F . In Example 1, this means that λ varies from 1 to some upper bound $\bar{\lambda}$. In Example 2, this means that $\lambda(H)$ varies similarly. In Example 3, this means that each type has a different $\lambda(\gamma_p)$ function. In the case of Example 3, include a single-crossing assumption: $\lambda(0) = 1$ for all types, and no two types otherwise intersect. How does this affect the producer's pricing problem?

The producer still only has two periods with which to price discriminate, and consumers are only willing to pay a premium due to λ in the early adoption stage. Thus, the producer faces a tradeoff between early adoption volume and early adoption price. It is immediate that the optimal early adoption pricing policy is some cutoff price p^* and associated cutoff type λ^* such that consumers with $\lambda \geq \lambda^*$ will early adopt and the rest will wait. In terms of equilibrium outcome, however, this is exactly equivalent to the two-type case where $\beta = 1 - F(\lambda^*)$ and (in Examples 1 and 2) $\lambda_e = \lambda^*$ or (in example 3) $\lambda_e(1 - F(\lambda^*)) = \lambda^*$.²⁸

Notice that this is true of any distribution of types (or in the case of Example 3, any that satisfies the single-crossing condition). It is for this reason that I chose not to impose a particular distribution (e.g. the uniform distribution) in the main body of the paper; doing so weakens, rather than strengthens, the results.

Editorial Policy without a Binding Contract

In the model, editorial policy is assumed to be a lifelong binding agreement such that the reviewer either complies and receives its lifetime discounted revenue stream R or it gets

²⁸Equilibrium payoffs are different; in the two-type case, the producer captures all of the surplus, but that is not the case here. However, equilibrium outcomes are equivalent.

nothing. This simplification allows for modeling an infinitely repeated process using just a single stage. Suppose, however, that the editorial policy agreement were not binding, the game were played every period, and the reviewer has the opportunity to “betray” the producer by deviating to some review policy other than the agreed upon one. Specifically, the reviewer will be tempted not to inflate a review in order to minimize reputational losses.

There are two possible scenarios: One where the producer randomizes in the editorial policy (or is at least aware of the randomization), and another where the reviewer does the randomization. In the former case, the producer can perfectly monitor the behavior of the reviewer. In the latter case, the producer’s monitoring is imperfect. In either case, there is a terrific amount of equilibrium multiplicity, but we can make inferences about the type of equilibrium being played with real-life observations.

When the bargaining relationship between the producer and the reviewer is highly asymmetric, we might expect the behavior of the reviewer to be largely cooperative due to the producer’s capacity for punishment. If this is the case, the profit maximizing approach for the producer is to extract exactly enough surplus from the reviewer that its long-term average payoff is zero (i.e. bind its IR constraint). In fact, this case yields an outcome identical to the one where the editorial policy is binding. Moreover, a lack of defection from reviewers is consistent with real-world observation; we see comparatively few instances of punishment relative to the immense number of reviews that exist. Thus, the assumptions in the model may be alternatively thought of as a case where policy is non-binding but the producer uses its bargaining power to shift equilibrium selection in its favor.

Why Even Have a Reviewer?

Given the construction of the model, a reader might wonder why a third-party reviewer is even necessary for this process – couldn’t the producer simply “own” a reviewer and have

full control over the transmitted review? The answer lies in how the optimal editorial policy in an inflation equilibrium is derived from the reviewer's IR constraint.

First, it is not clear if a "first-party reviewer" should suffer the same reputational costs that the third-party reviewer does in the model. Firms often do not even try to present themselves as impartial evaluators of their own products, making public statements about their products mere cheap talk or puffery. This is precisely why firms often appeal to third parties even within their own advertisements ("four out of five dentists agree!"). Since reputational costs are the only force that keeps the optimal editorial policy from being $\eta = 1$ in an inflation equilibrium, we would expect zero information transmission in this case.

Even if we suppose that first-party reviewers suffer reputational costs in exactly the same way that third-party reviewers do, their optimal editorial policy will still be different in an inflation equilibrium. Suppose a producer also owns and operates the reviewer. Then, from Proposition 2, its total profit in an inflation equilibrium is:

$$\begin{aligned} \pi = \beta\theta H(\lambda_e + (1 - \lambda_e)\nu) + \beta(1 - \theta)L(1 + (\lambda_e - 1)\eta) + (1 - \beta)(\theta H + (1 - \theta)L) \\ + R - [(1 - \theta)\eta + \theta\nu]\phi \end{aligned}$$

It is immediate that $\frac{\partial\pi}{\partial\nu} < 0$ and therefore $\nu^* = 0$. When can then write π as a function of η :

$$\pi(\eta) = \beta\theta H\lambda_e + \beta(1 - \theta)L(1 + (\lambda_e - 1)\eta) + (1 - \beta)(\theta H + (1 - \theta)L) + R - [(1 - \theta)\eta + \theta\nu]\phi \quad (48)$$

$$\Rightarrow \frac{\partial\pi}{\partial\eta} = \beta(1 - \theta)L(\lambda_e - 1) - (1 - \theta) \quad (49)$$

This maximization problem has a bang-bang solution; with the exception of the knife-edge case, the derivative of the firm's profit is either strictly positive or strictly negative, so η^* will either be equal to 0 or 1. The case where $\eta^* = 0$ is the case where the costs of misrepresentation are very high, resulting in full self-disclosure (for example: disclaimers on pharmaceuticals to prevent lawsuits). In the alternative case, however, the review is once again merely cheap talk, and if consumers have a preference for partially informative signals over uninformative ones, we would expect third-party reviewers to eventually force first-party reviewers out of the market. This is precisely what occurred in the video game industry; Microsoft, Sony, and Nintendo all once published their own video game magazines, but all three are now either defunct or have been outsourced to third parties.

References

Basuroy, S., S. Chatterjee, and S.A. Ravid (2003). “How Critical are Critical Reviews? The Box Office Effects of Film Critics, Star Power, and Budgets,” *Journal of Marketing* 67(4), 103-117.

Basuroy, S., and S.A. Ravid (2013). “How Relevant Are Experts in the Internet Age? Evidence from the Motion Pictures Industry,” *Working Paper*.

Blasco, A., F. Sobbrío, and P. Pin (2015). “Paying Positive to Go Negative: Advertisers’ Competition and Media Reports,” *Working Paper*, available at SSRN.

Camara, F., and N. Dupuis (2014). “Structural Estimation of Expert Strategic Bias: The Case of Movie Reviewers,” *Working Paper*, available at SSRN.

Caves, R. (2000). *Creative Industries: Contracts Between Art and Commerce*. Cambridge, Mass.: Harvard University Press.

Cox, J., and D. Kaimann (2014). “The Signaling Effect of Critics: Do Professionals Outweigh Word-of-Mouth? Evidence from the Video Game Industry,” *Working Paper*.

Crawford, V., and J. Sobel (1982). “Strategic Information Transmission,” *Econometrica* 50(6), 1431-1452.

Dewenter, R., and U. Heimeshoff (2014). “Media Bias and Advertising: Evidence from a German Car Magazine” *Review of Economics* 65(1), 77-94.

Dewenter, R., and U. Heimeshoff (2015). “Do Expert Reviews Really Drive Demand? Evidence from a German Car Magazine.” *Applied Economics Letters* 22(14), 1150-1153.

Dobrescu, L., M. Luca, and A. Motta (2013). “What Makes a Critic Tick? Connected Authors and the Determinants of Book Reviews,” *Journal of Economic Behavior and Organization* 96, 85-103.

- Dranove, D., and G. Jin (2010). "Quality Disclosure and Certification: Theory and Practice," *Journal of Economic Literature* 48(4), 935-963.
- Elberse, A., and J. Eliashberg (2003). "Demand and Supply Dynamics Behavior for Sequentially Released Products in International Markets: The Case of Motion Pictures," *Marketing Science* 22(3), 329-354.
- Ellman, M., and F. Germano (2009). "What Do Papers Sell? A Model of Advertising and Media Bias," *Economic Journal* 119, 680-704.
- Friberg, R., and E. Gronqvist (2012). "Do Expert Reviews Affect the Demand for Wine?" *American Economic Journal: Applied Economics* 4(1), 193-211.
- Gentzkow, M., and E. Kamenica (2011). "Bayesian Persuasion," *American Economic Review* 101(6), 2590-2615.
- Goh, J., and L. Ederington (1993). "Is a Bond Rating Downgrade Bad News, Good News, or No News for Stockholders?" *Journal of Finance* 48(5), 2001-2008.
- Greco, A., J. Milliot, and R. Wharton (1997). *The Book Publishing Industry*. Allyn Bacon, Needham Heights: MA.
- Hamami, T. (2016). "Expert Product Reviews and Conflict of Interest," *Working Paper*, available at SSRN.
- Hennig-Thurau, T., M. Houston, and G. Walsh (2006). "The Differing Roles of Success Drivers Across Sequential Channels: An Application to the Motion Picture Industry." *Journal of the Academy of Marketing Science* 34(4), 559-575.
- Hernandez, M., and H. Vicdan (2014). "Modeling word of mouth vs. media influence on video game preorder decisions: A qualitative approach," *Journal of Retailing and Consumer Services* 21, p. 401-406.

- Hilger, J., G. Rafert, and S. Villas-Boas (2011). "Expert Opinion and the Demand for Experience Goods: An Experimental Approach," *Review of Economics and Statistics* 93(4), 1289-1296.
- Hu, N., P. Pavlou, and J. Zhang (2009). "Overcoming the J-Shaped Distribution of Product Reviews," *Communications of the ACM* 52(10), 144-147.
- Kartik, N. (2009). "Strategic Communication with Lying Costs," *Review of Economic Studies* 76(4), 1359-1395.
- Kolotinin, A. (2015). "Optimal Information Disclosure: Quantity vs. Quality," *Working Paper*, available at SSRN.
- Peltzman, S. (1976) . "Toward a More General Theory of Regulation" *Journal of Law and Economics* 19(2), 211-240.
- Reinstein, D., and C. Snyder (2005). "The Influence of Expert Reviews on Consumer Demand for Experience Goods: A Case Study of Movie Critics," *Journal of Industrial Economics* 53(1), 27-51.
- Reuter, J. (2009). "Does Advertising Bias Product Reviews? An Analysis of Wine Ratings." *Journal of Wine Economics* 4(2), 125-151.
- Rinallo, D., and S. Basuroy (2009). "Does Advertising Spending Influence Media Coverage of the Advertiser?" *Journal of Marketing* 73(6), 33-46.
- Stigler, G. (1971). "The Theory of Economic Regulation" *Bell Journal of Economics and Management Science* 2(1), 3-21.
- Wingfield, N. (2007). "High Scores Matter to Game Makers, Too," *The Wall Street Journal* September 2007. Available at <http://www.wsj.com/articles/SB119024844874433247>.
- Zhu, F., and X. Zhang (2010). "Impact of Online Consumer Reviews on Sales: The Moderating Role of Product and Consumer Characteristics." *Journal of Marketing* 74(2), 133-148.