TOPICS IN INTELLECTUAL PROPERTY

Peter Adam Malaspina

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics.

Chapel Hill

2011

Approved by,

Peter Norman

Gary Biglaiser

Helen Tauchen

Sergio Parreiras

Brian McManus

ABSTRACT

Peter Malaspina: Topics in Intellectual Property
(Under the instruction of Dr. Peter Norman)

The following chapters cover theoretical and empirical investigations into issues of intellectual property.

The first chapter examines predatory behavior in patent litigation. I develop a signaling model where the timing of litigation against an initial act of infringement transmits noisy signals about a patent holder's private expectations of litigation awards. I find a non-monotonic relationship between a patent holder's award expectation and the timing of litigation. Patent holders with sufficiently high award expectations will exhibit predatory behavior by delaying litigation to lure infringers. At the same time, patent holders with moderate award expectations will litigate quickly to deter additional acts of infringement.

The second chapter (co-authored with Dr. David Molin) examines the economic consequences of the reduction in the fixed costs of distribution due to the emergence of Internet and information technology. Others have examined the effects of increasing product variety on consumer welfare, in bookselling and in other markets. We identify a set of books sold on Amazon.com for which the author's revenue can be easily computed. We use this data to parameterize a model of author entry and then estimate the welfare gains of authors and consumers from Print-on-Demand (POD) technology.

Table of Contents

List of Tables

Table

List of Figures

Figure

CHAPTER 1

PREDATORY DELAYS IN PATENT LITIGATION

1.1 Introduction

In 2007 a jury awarded the plaintiff Alcatel-Lucent $1.5 billion for the violation of two patents by the defendant Microsoft, who allegedly used unlicensed audio technology in their Windows operating systems beginning in 1997. After siding with Alcatel-Lucent, the court based their award on 0.5% of the retail value of all the computers sold with the unlicensed technology. Furthermore, Alcatel-Lucent waited until after the release of Windows XP in 2001 to pursue litigation in 2002, well after the first unlicensed use of the patented technology in 1997. So the question remains: why did Alcatel-Lucent wait so long to initiate litigation?

One theory is that Alcatel-Lucent had an incentive to delay litigation until after Microsoft committed a subsequent act of infringement with the release of Windows XP. Alcatel-Lucent may have foreseen the potential for a larger award in the future by noting that sales were accumulating during the delay, and that these sales would probably be used to calculate awards. Therefore, if Alcatel-Lucent had initiated litigation based on earlier versions of Windows, Microsoft may have avoided incorporating the patented technology into future products altogether, and thus eliminated the potential for Alcatel-Lucent to extract any royalty for subsequent use.

The example above suggests that a patent holder's award expectation may affect the timing of litigation when there is a potential for additional infringement in the future. In this paper, I model how this process works and analyze its impact on the behavior of patent holders and infringers. To this end, I develop a signaling model where the timing of litigation against an act of infringement transmits information about a patent holder's private expectations of litigation awards.

In the analysis of my model I find many cases where patent holder behavior is impacted by two reputation effects: deterrence and luring. Patent holders with sufficiently high award expectations will utilize luring by delaying litigation to encourage additional acts of infringement.[1] Patent holders with relatively moderate award expectations will utilize deterrence by litigating quickly before additional acts of infringement can be committed. Last, patent holders with the lowest award expectations will create no litigation. Therefore, the probability of litigation preempting any additional infringement is increasing in the patent holder's award expectation up to a point, after which the probability of preemptive litigation is decreasing until it reaches zero.[2]

The incentives for predatory patent holder behavior exist under the following conditions: First, at least one patent holder type must expect litigation in face of infringement to return greater expected profits than the best alternative in the absence of infringement. Second, the infringer's equilibrium entry decision must be conditioned on their observations of any previous litigation.[3] Together, these conditions will imply that there exists at least one type of patent holder that has an incentive to delay litigation for the sake of predation.

I show that there exists a unique perfect Bayesian equilibrium for each parameter setting of the model. The existence property of the model, combined with the intuitive nature of resulting equilibrium, demonstrates the model's ability to provide insightful results for a wide range of settings. Furthermore, the uniqueness property shows that my conclusions are unbiased with respect to the equilibrium I choose to analyze. Uniqueness is a useful property for models with reputation, which usually result in multiple equilibria,

[1] Litigation Awards can be up to three times damages plus legal costs. Furthermore, since awards for patent infringement cannot fall below the courts definition of a "reasonable royalty", a patent holder who is confident in the validly of his patent, as well as his ability to prove infringement, should view an expected licensing fee as an approximate baseline for awards.

[2] In our model, a patent holder that desires deterrence is analogous to a "weak" incumbent acting "tough" in the classic entrant incumbent game modeled in Kreps and Wilson(1982). However the potential incentives for luring provide a new twist on an old game. That is, we also see tough types acting weak.

[3] In some cases the infringer will ignore the first round signal. For example, when no patent holder type has a positive award expectation, the infringer will commit additional infringement regardless of the observed signal, because no feasible posterior beliefs could deter an infringer from entry. In this case the infringer's equilibrium entry decision is unconditional.

of which at least some are unaffected by reputation.

The predatory behavior found in my model is not necessarily prohibited under current U.S. patent law. The current statute of limitations requires that litigation be brought within 6 years of an initial act of infringement.[4] In many cases, opportunities for the infringer to commit additional acts of infringement exist well within this time frame.[5] However, I show that policy makers can affect the magnitude of reputational distortions without changing the statute of limitations. To this end, I demonstrate how the magnitude of reputation's effect on behavior is linked to the prior distribution of expected awards, noting that policy makers have some control over these distributions.

My research is related to other work that focuses on the strategic behavior of patent holders in the face of potential infringement. Crampes and Langinier (2002) model patent holders who must expend effort to monitor for infringement. They calculate the optimal effort level, and find that increasing awards may not necessarily decrease the probability of infringement. I include a similar notion of patent holder effort. Michael J. Meurer (1989) models the nature of settlement offers to a potential infringer. Meurer analyzes how a patent holder's private expectations of their patent's validity can lead to a particular settlement, litigation, or inaction outcome. My model also assumes that patent holders have private information about their patent. However, I expand on Meurer's notion of private information by including any element of private information that affects expected awards. Certainly, information about validity will affect a patent holder's expected award. Therefore, my interpretation of private information includes, but is not limited to, issues of validity. Neither Crampes and Langinier nor Meurer, considers a repeated interaction between a patent holder and a sequence of infringers. Therefore, issues concerning information revelation and reputation effects are not present.

The closest to modeling reputation effects in patent litigation is Choi (1998), who examines the implications of information transmission on the strategic behavior of an endogenous sequence of infringers. While there are some similarities in our models, there are also some important differences: In his model all patent holders are ex-ante

[4] Pincus(1991)

[5] Any future production or sales of an infringing product after the initial act of infringement could constitute an additional act of infringement in the model.

identical. Therefore, it is impossible to account for the different reputational incentives of patent holders who differ in their private expectations of litigation awards. Choi also assumes that information is transmitted the moment litigation begins; whereas in my model, litigation is ongoing, such that it does not fully reveal private information to other potential infringers. The distinctions lead us to different conclusions about why patent holders might want to delay litigation. Whereas, Choi finds that patent holders will delay litigation because they are afraid of revealing weakness, I find that patent holders may delay litigation to avoid revealing strength.

The rest of this paper is presented in the following order: The second section presents the model preliminaries, and defines equilibrium strategies and beliefs. The third section presents the resulting equilibrium properties of the model. The fourth section discusses the properties of the model as they relate to reputation, potential policy, and signaling games in general. The fifth section states my conclusions. The sixth section is an appendix containing the proofs of all Lemmas and Propositions.

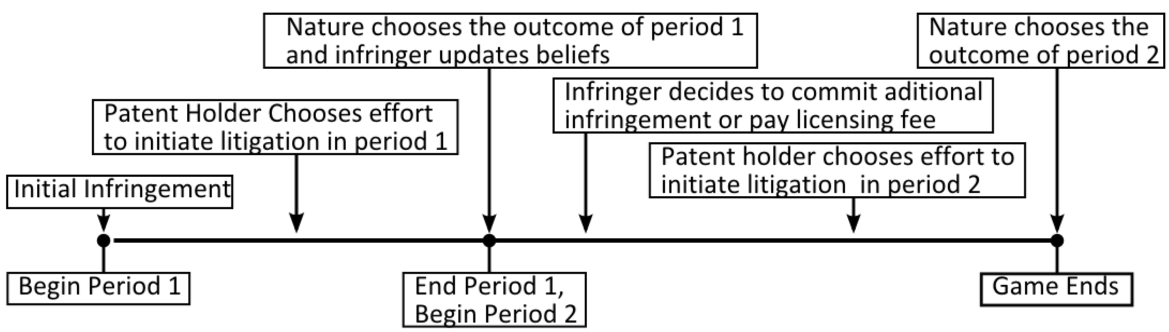

Figure 1: A timeline of the game

1.2 Model

Consider a two period model where a patent holder faces off against an initial act of infringement committed at the beginning of period 1.[6] Furthermore, the patent holder foresees the potential of an additional act of infringement in period 2.[7] Assume patent holders have private information about their expected awards from any litigation. In the resulting signaling game, signals about the patent holder's private information will be transmitted by the outcome of period 1 to the potential infringer before they are called to make an entry decision in period 2. The timing of this game is summarized in figure 1.

Assume patent holders vary in their ex-ante expected awards from any ongoing litigation that has been successfully initiated, such that awards are correlated with private information θ for each type $\theta \in \Theta \subset R$.[8]

Furthermore, let $\phi_0(\cdot)$ denote the infringer's infringer's prior beliefs about the patent holder's private information in period 1. Such that $\phi_0(\cdot)$ assigns each type θ a prior probability density, denoted as $\phi_0(\theta)$.

Let $\gamma > 0$ be an exogenous scalar that captures the relative size of the potential act of infringement to the initial act, such that

1. If litigation is initiated successfully against either act *separately*, the patent holder would expect award θ for the initial act of infringement, and $\theta\gamma$ for the potential act of infringement.

2. If litigation is initiated successfully against both acts *jointly*, the expected award becomes $(\theta + \theta\gamma)$.

The game begins when a patent holder faces an initial act of infringement in period 1.

[6] First period infringement is assumed for simplicity. Endogenous first period entry will not change the general intuition of our results.

[7] This model depicts what happens when the potential infringer is the same as the initial infringer (See the Microsoft example in the introduction). If instead I want to analyze what happens when the potential infringer is another firm, I can make minor modifications to the model that will not affect the intuition of the results.

[8] I allow for types $\theta < 0$ for generality. $\theta < 0$ might be representative of a patent holder with invalid property rights who might expect a countersuit.

In order to successfully initiate litigation, the patent holder must expend effort denoted e_1. Based on the patent holder's effort, nature determines the the probability $p(e_1)$ that the patent holder successfully finds arguments sufficient to bring the case to court in period 1. I assume $e_1 \in R_+$, and $p(e_1)$ is twice continuously differentiable such that $p^{'}(e_1) > 0$, $p^{''}(e_1) < 0$, and $p(0) = 0$.[9] Therefore, based on his effort, the patent holder creates litigation in period 1 with probability $p(e_1)$, with complementary probability $(1 - p(e_1))$ no litigation exists at the beginning of period 2.

Any litigation initiated in period 1 is on-going at the beginning of period two. Therefore, the only signal observed by the infringer in period 2 is the history $h_1 \in \{L, \not{L}\}$, where $\{L\}$ denotes the **existence,** and $\{\not{L}\}$ denotes the **absence,** of ongoing litigation from period 1.[10]

Therefore, I denote the infringer's period 2 entry action as being conditional on observing h_1. Let $\sigma_2(h_1) \in [0,1]$ denote the conditional entry probability, as mixed strategies over $\{Out, In\}$. If the infringer decides not to commit additional acts of infringement, they pay an exogenous licensing fee F to the patent holder for additional use.[11] On the other hand, if they commit additional acts of infringement, they avoid the licensing fee, but may be subject to litigation in period 2.

After the potential infringer's decision in period two, there are four possible histories facing a patent holder in period 2, stemming from the two possible histories of period 1 $\{L, \not{L}\}$, and the two possible actions of the infringer in period 2 $\{Out, In\}$ (See Figure 2). Therefore, the expected award for successfully initiated litigation in period 2 will be conditional on a patent holder's type (θ) and the set of observable histories $h_2 \in \{h_1, a_2(h_1)\}$. Here I summarize the four possible histories and provide the corresponding

[9] It is important to keep in mind that failure to initiate litigation does not mean that a patent holder loses in court. In this model, failure to initiate litigation in period 1 is interpreted as a failure to find arguments sufficient to bring the case to court in period 1.

[10] I defend the assumption about the observable history as follows: litigation is usually a long process. Protracted legal battles can delay information transmission beyond the onset of a trial. For example, in the case of Polaroid V. Kodak, it took 14 years for the patent infringement suit to be resolved. Therefore, the infringer will likely have other opportunities to commit additional acts of infringement before they can observe the final award.

[11] The implications of an endogenous licensing fee are examined in the discussion section. I show that with common sense assumptions, allowing for the endogenous licensing will not change the intuition of the results.

expected player profits.

Let $\Pi^{PH}(h_2, \theta)$ denote a type θ patent holder's expected profits upon reaching history h_2. Furthermore, let $\Pi^{IN}(h_2 \mid \theta)$ denote the infringer's expected profits upon reaching history h_2, conditional on facing a patent holder of type θ.[12] The four potential histories facing the patent holder in period 2 are defined as follows:

History $h_2 = \{L, Out\}$: The infringer decides not to commit an additional act of infringement in the presence of ongoing litigation from period 1. The infringer pays an exogenous licensing fee (F) to the patent holder for additional use, and awaits the outcome of the ongoing litigation from round 1.

$$\Pi^{PH}(L, Out, \theta) = \theta + F - e_1 \tag{1}$$

$$\Pi^{IN}(L, Out \mid \theta) = -\theta - F \tag{2}$$

[12] Note that at this point in the game, the infringer will not necessarily know what type of patent holder it faces. Thus, the infringer's equilibrium action will depend on its beliefs. Beliefs are detailed in section 2.3.

History $h_2 = \{L, In\}$ The infringer commits an additional act of infringement in the presence of ongoing litigation from period 1. The infringer avoids paying the licensing fee. However, the patent holder can then initiate litigation against the additional act of infringement separately, with certainty, and without any additional cost of effort.

$$\Pi^{PH}(L, In, \theta) = \theta + \theta\gamma - e_1 \tag{3}$$

$$\Pi^{IN}(L, In \mid \theta) = -\theta - \theta\gamma \tag{4}$$

History $h_2 = \{\not{L}, In\}$: the infringer commits an additional act of infringement in the absence of ongoing litigation from period 1. The infringer avoids paying the licensing fee. However, the patent holder can expend effort (e_2) to try and initiate litigation against both the initial and additional acts of infringement jointly.

$$\Pi^{PH}(\not{L}, In, \theta) = (\theta\gamma + \theta)\, p\left(e_2^*(\theta \mid In)\right) - e_2^*(\theta \mid In) - e_1 \tag{5}$$

$$\Pi^{IN}(\not{L}, In \mid \theta) = -(\theta\gamma + \theta)\, p\left(e_2^*(\theta \mid In)\right), \tag{6}$$

$$where\ e_2^*(\theta \mid In) = \underset{e_2 \in R_+}{\operatorname{argmax}} \left\{(\theta\gamma + \theta)\, p(e_2) - e_2\right\}.$$

History $h_2 = \{\not{L}, Out\}$: The infringer decides not to commit an additional act of infringement in the absence of litigation from period 1. The infringer pays a licensing fee for additional use of the patent, but the patent holder can still expend effort (e_2) to litigate the initial act of infringement from period one.

$$\Pi^{PH}(\not{L}, Out, \theta) = F + \theta p\left(e_2^*(\theta \mid Out)\right) - e_2^*(\theta \mid Out) - e_1 \tag{7}$$

$$\Pi^{IN}(\not{L}, Out \mid \theta) = -F - \theta p\left(e_2^*(\theta \mid Out)\right), \tag{8}$$

$$where\ e_2^*(\theta \mid Out) = \underset{e_2 \in R_+}{\operatorname{argmax}} \left\{\theta p(e_2) - e_2\right\}.$$

All payoffs are realized at the end of the second period. I now define a reduced form of a perfect Bayesian equilibrium of my model.

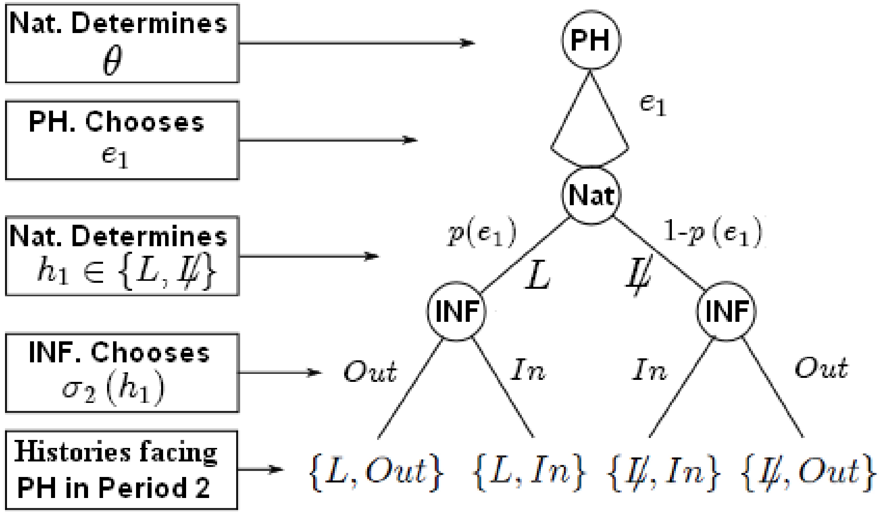

Figure 2: An overview of the game.

The reduced form[13] of a **Perfect Bayesian Equilibrium** will consist of

1. Each patent holder type's effort levels: $\{e_1^*(\theta)\}\ \forall \theta \in \Theta$
2. The infringer's entry probabilities: $\sigma_2^*(L), \sigma_2^*(\not{L})$.
3. The infringer's updated beliefs: $\{\phi(\cdot \mid L), \phi(\cdot \mid \not{L})\}$.

Such that:

- Given $\sigma_2^*(L)$ and $\sigma_2^*(\not{L})$, $e_1^*(\theta)$ maximizes the expected profits of each type θ.
- Given $\phi(\cdot \mid h_1)$, $\sigma_2^*(h_1)$ is profit maximizing for the infringer for all $h_1 \in \{L, \not{L}\}$.
- The infringer's updated beliefs $\phi(\cdot \mid h_1)$ are consistent with $\{e_1^*(\theta)\ \forall\ \theta \in \Theta\}$ for all $h_1 \in \{L, \not{L}\}$.[14]

[13] I omit $e_2^*(\theta \mid Out), e_2^*(\theta \mid In)$ from the reduced form because they are unaffected by beliefs. See (7) and (5).

[14] Perfection is satisfied when there is at least one type of patent holder with positive litigation expectations because both possible signals will occur with positive probability in equilibrium. In this case, there will be no off equilibrium histories to generate off equilibrium beliefs. Furthermore, when no patent holder type has a positive expectation of litigation awards, I can assume that the infringer always enters in period 2 because no posterior beliefs over the support of initial beliefs could deter additional infringement.

Solving for Equilibrium In the signaling game I model the infringer observes a noisy signal of the patent holder's award expectation based on their understanding of each patent holder type's first period optimization problem. Therefore, while an infringer does not observe the patent holder's first period effort, they can still update their beliefs in a manner consistent with the equilibrium effort of the patent holder. This consistency results in a useful property: that a change in equilibrium entry probability will always affect the equilibrium effort level of at least one patent holder type. Therefore any potential change in equilibrium infringer entry probability will change their expected profits. Therefore, the resulting PBE of the game will always be unique. Thus I prove that my findings about predatory behavior generated by the model are general, and not the result of a selective choice of the equilibria to analyze.

The general procedure for identifying equilibria is guess and verify, which I briefly summarize as follows:

1. Guess a set of infringer best responses

$$\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} \tag{9}$$

2. Given $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\}$, calculate each type patent holder's optimal effort.

$$\{\hat{e}_1^*(\theta)\,\forall\,\theta \in \Theta\} = \{e_1^*(\theta \mid \hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L}))\,\forall\,\theta \in \Theta\} \tag{10}$$

3. Given $\{\hat{e}_1^*(\theta)\,\forall\,\theta \in \Theta\}$, calculate the infringer's updated beliefs

$$\hat{\phi}(\theta \mid L) = \phi(\theta \mid L, \hat{e}_1^*(\theta)\,\forall\,\theta \in \Theta) \tag{11}$$

$$\hat{\phi}(\theta \mid \not{L}) = \phi(\theta \mid \not{L}, \hat{e}_1^*(\theta)\,\forall\,\theta \in \Theta) \tag{12}$$

4. Using updated beliefs, verify that best responses are indeed best responses

$$\hat{\sigma}_2^*(L) = \sigma_2^*(L \mid \{\hat{e}_1^*(\theta)\,\forall\,\theta \in \Theta\}) \tag{13}$$

$$\hat{\sigma}_2^*(\not{L}) = \sigma_2^*(\not{L} \mid \{\hat{e}_1^*(\theta)\,\forall\,\theta \in \Theta\}) \tag{14}$$

In the following I sections I will describe the relationships between optimal patent holder effort, the infringer's beliefs, and the infringer's best responses.

Updating Beliefs

Before making an entry decision in period 2, the infringer will update its beliefs about the distribution of patent holder types based on the observable outcome of period 1, its prior beliefs $\phi_0(\cdot)$, and the set of equilibrium effort levels for each patent holder type $\{e_1^*(\theta)\,\forall\,\theta\in\Theta\}$. Note that effort is not observable; however, within any equilibrium, I show below that each patent holder type has a unique optimal amount of effort. Therefore, $\{e_1^*(\theta)\,\forall\,\theta\in\Theta\}$ can be used to update beliefs in equilibrium, because the infringer will know the optimization problem faced by each patent holder type. Here I define the infringer's updated beliefs as the set of density functions $\{\phi(\cdot \mid h_1)\ \forall h_1\in\{L,\not{L}\}\}$.

Definition 1 *The infringer's updated beliefs, is the mapping*

$$\phi:\{\Theta,h_1\}\longrightarrow[0,1] \tag{15}$$

such that, upon observing h_1

$$\phi(\theta\mid L) = \frac{p(e_1^*(\theta))\,\phi_0(\theta)}{\int\limits_{\theta\in\Theta} p(e_1^*(\theta))\,\phi_0(\theta)\,d\theta} \tag{16}$$

$$\phi(\theta\mid \not{L}) = \frac{(1-p(e_1^*(\theta)))\,\phi_0(\theta)}{\int\limits_{\theta\in\Theta}(1-p(e_1^*(\theta)))\,\phi_0(\theta)\,d\theta} \tag{17}$$

where $p(e_1^*(\theta))$ *is the probability a type* θ *patent holder creates litigation in period 1 given its equilibrium first period effort* $e_1^*(\theta)$, *and* $\phi_0(\theta)$ *is the prior density of a type* θ *in nature.*

The Infringer's Entry Decision

There are only two subgames facing the potential infringer stemming from the two observable outcomes in period one: the existence of litigation, or lack thereof, denoted respectively as the history $h_1\in\{L,\not{L}\}$. This outcome will allow the infringer to update their beliefs from $\phi_0(\cdot)$ to $\phi(\cdot \mid h_1)$. Therefore the potential infringer's expected payoffs from infringement become conditional on the outcome of round 1, and their entry decision.

To see if the infringer is maximizing profits by choosing $\{In\}$ in period 2, I must compare it to their reservation value from choosing $\{Out\}$. Note, that an infringer's expected profits from entry and their expected reservation value is conditional on the observed the first period outcome. Here I define their net expected payoffs from entry under each observable first period history. Notice that the infringer's expected profits consist of its expected profits conditioned on patent holder type within the corresponding second period histories (see (2), (4), (6), and (8)), and the densities assigned by their updated beliefs.

$$\Pi^{IN}(h_1, In) = \int_{\theta\in\Theta} \left[\Pi^{IN}(h_1, In \mid \theta) - \Pi^{IN}(h_1, Out \mid \theta)\right] \phi(\theta \mid h_1) d\theta \tag{18}$$

I now define an infringer best response for each observed history form period 1, as a function of the net expected payoffs from entry.

Definition 2 *The best response for the infringer, given an observed first period history* h_1*, is the mapping*

$$\sigma_2^* : \{L, \not{L}\} \longrightarrow [0,1] \tag{19}$$

such that

$$\begin{aligned} \sigma_2^*(h_1) &= 1, && \text{if } \Pi^{IN}(h_1, In) > 0 \\ \sigma_2^*(h_1) &\in (0,1), && \text{if } \Pi^{IN}(h_1, In) = 0 \\ \sigma_2^*(h_1) &= 0, && \text{if } \Pi^{IN}(h_1, In) < 0 \end{aligned} \tag{20}$$

where $\sigma_2^*(h_1)$ *is their probability of committing the additional act of infringement conditional on observing the first period history* h_1.

The Patent Holder's Optimization Problem

The patent holder's optimization problem depends on the equilibrium best responses of the infringer to the first period outcome. Therefore, any set of best responses $\{\sigma_2^*(L), \sigma_2^*(\not{L})\}$ will lead to a unique choice of effort$(e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L})))$ for each patent holder type $\theta \in \Theta$. Hence for each patent holder type $\theta \in \Theta$,

$$e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L})) = \arg\max_{e_1 \in R_+} \left\{ \begin{array}{l} p(e_1)\sigma_2^*(L)\,\Pi^{PH}(L, In, \theta) + \\ (1 - p(e_1))\,\sigma_2^*(\not{L})\,\Pi^{PH}(\not{L}, In, \theta) + \\ p(e_1)\,(1 - \sigma_2^*(L))\,\Pi^{PH}(L, Out, \theta) + \\ (1 - p(e_1))\,(1 - \sigma_2^*(\not{L}))\,\Pi^{PH}(\not{L}, Out, \theta) \end{array} \right\}.^{15} \quad (21)$$

As I see from the patent holder's profit maximization problem above, the patent holder considers the likelihood of reaching any history given its choice of effort in period 1 and the set of equilibrium infringer best responses $\{\sigma_2^*(L), \sigma_2^*(\not{L})\}$. A patent holder must also consider its expected profits given that it reaches a particular history, given in (1) (3) (5) and (7).[16]

The patent holder's optimal first period effort is based on the following first order condition. Here I present this condition with respect to e_1.

$$p'(e_1)\,\psi(\theta) \leq 1, \quad (22)$$

where $\psi(\theta)$ is defined as

$$\psi(\theta) = \left[\begin{array}{l} \Pi^{PH}(L, Out, \theta) - \Pi^{PH}(\not{L}, Out, \theta) + \\ \sigma_2^*(L)\left[\Pi^{PH}(L, In, \theta) - \Pi^{PH}(L, Out, \theta)\right] + \\ \sigma_2^*(\not{L})\left[\Pi^{PH}(\not{L}, Out, \theta) - \Pi^{PH}(\not{L}, In, \theta)\right] \end{array} \right], \quad (23)$$

[16] Note that the patent holders expected profits, given that they reach a particular subgame, are ex-ante predetermined (as illustrated in descriptions of the period 2 subgames above). Therefore I do not restate these profit maximizing period two choices of effort here because they are static with respect to any potential equilibrium.

and the corresponding second order condition

$$p''(e_1)\psi(\theta) < 0. \tag{24}$$

Notice that a given set of infringer best responses $\{\sigma_2^*(L), \sigma_2^*(\not L)\}$ is sufficient to identify the optimal first period effort for each patent holder type. Notice that the first order condition holds with equality for a given set of best responses $\{\sigma_2^*(L), \sigma_2^*(\not L)\}$ whenever $\psi(\theta) > 0$, because $p(e_1)$ is concave and twice continuously differentiable, and thus the second order condition is also satisfied. Furthermore, for all types θ such that $\psi(\theta) \leq 0$, I know $e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not L)) = 0$ because the first order condition will never bind. Therefore, for a given set of best responses each patent holder type has only one level of effort that maximizes its expected profits. This property is essential for showing that any equilibrium is unique.

1.3 Results

In this section, I show that any parameterization of the model, with any distribution of types $\Theta \subset R$, and any twice continuously differentiable function $p(\cdot)$ such that $p'(e) > 0$, $p''(e) < 0$, $p(0) = 0$, has a unique Perfect Bayesian equilibrium.

First, I prove in Lemma 1 that the patent holder's choice of first period effort is a continuously differentiable function of the patent holder's type θ, and the infringer's conditional equilibrium entry probabilities: $\sigma_2^*(L)$ and $\sigma_2^*(\not L)$. This property is reliant on allowing the patent holder's choice of e_1 to be continuous. This will allow me to use the appropriate derivatives in my proofs of the subsequent Lemmas and Propositions.

Lemma 1 *For all $\theta \in R$, $e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not L))$ is continuously differentiable in $\sigma_2^*(L)$, $\sigma_2^*(\not L)$, and θ.*

Next, lemmas 2 and 3 show how a change in the infringer's equilibrium conditional entry probability affects the equilibrium relative first period effort of each patent holder type. These lemmas will help me prove lemmas 4 and 5, and assist me in my analysis patent holder effort in the discussion section.

Lemma 2 *As $\sigma_2^* (\not{L})$ increases, the change in first period patent holder effort is a weakly decreasing function of patent holder type θ, such that*

$$\frac{\partial^2 e_1^* \left(\theta \mid \sigma_2^* (L), \sigma_2^* (\not{L})\right)}{\partial \sigma_2^* (\not{L}) \, \partial \theta} \leq 0, \ \forall \theta \tag{25}$$

Lemma 3 *As $\sigma_2^* (L)$ increases, the change in equilibrium first period effort is a weakly increasing function of type θ.*

$$\frac{\partial^2 e_1^* \left(\theta \mid \sigma_2^* (L), \sigma_2^* (\not{L})\right)}{\partial \sigma_2^* (L) \, \partial \theta} \geq 0, \ \forall \theta \tag{26}$$

Next, in lemmas 4 and 5 I show how a change in the infringer's equilibrium conditional entry probability (either $\sigma_2^* (L)$ or $\sigma_2^* (\not{L})$) affects its expected profits from entry after any first period history (h_1), via its affect on equilibrium patent holder effort shown in Lemmas 2 and 3. Lemma's 4, and 5 are necessary conditions to prove uniqueness and existence of equilibria in my model.

Lemma 4 *The infringer's expected profits from entry, after observing the existence of litigation in period 1 $(h_1 = L)$, is*
(4.1) *monotonically increasing in $\sigma_2^* (\not{L})$*
(4.2) *monotonically decreasing in $\sigma_2^* (L)$*

Lemma 5 *The infringer's expected profits from entry after observing the absence of litigation $(h_1 = \not{L})$ in period 1, is*
(5.1) *monotonically decreasing in $\sigma_2^* (\not{L})$*
(5.2) *monotonically increasing in $\sigma_2^* (L)$.*

Next, Lemma 6 will show how the probability of the infringer committing additional acts of infringement must be greater under the absence of litigation than it is under the existence of litigation. This assures me that in all equilibria the existence of litigation will cause more deterrence than the absence of litigation.

Lemma 6 *In any Bayesian equilibrium $\sigma_2^* (L) > 0$ only if $\sigma_2^* (\not{L}) = 1$.*

I now present Proposition 1, which summarizes the existence and uniqueness properties of equilibria in the model. It states that for any setting of the model there will exist one, and only one equilibrium.

Proposition 1 *For any setting of the parameters and any continuous function $p(\cdot)$ under my assumptions, and for any possible prior beliefs $\phi_0(\cdot)$ and any type space $\Theta \subset R$, there exists a unique perfect Bayesian equilibrium.*

The existence property of the model is proven by construction in the appendix. The proof of existence is outlined as follows: I begin by showing the procedure for checking for equilibrium in pure strategy infringer best responses using the method outlined in (14)-(19).

$$\{\sigma_2^*(L), \sigma_2^*(\not{L})\} \in \{\{Out, Out\}, \{Out, In\}, \{In, In\}\},$$

When there is no equilibrium in pure strategies, I use the following Lemmas, to show that there must be an equilibrium in mixed entry strategies. In other words, if

$$\{\sigma_2^*(L), \sigma_2^*(\not{L})\} \notin \{\{Out, Out\}, \{Out, In\}, \{In, In\}\} \tag{27}$$

then

$$\{\sigma_2^*(L), \sigma_2^*(\not{L})\} \in \{Out, \sigma_2^*\} \text{ or } \{\sigma_2^*, In\} \tag{28}$$

where $\sigma_2^* \in (0,1)$.

The uniqueness property of equilibrium in my model is derived from Lemmas 4 and 5; because even when the infringer is indifferent to entry under a certain history, and therefore playing a equilibrium mixed entry strategy $\sigma_2^*(h_1)$, deviating and playing $\tilde{\sigma}_2^*(h_1) \neq \sigma_2^*(h_1)$ will disrupt the equilibria by changing the equilibrium first period patent holder efforts, which will in turn affect the infringer's expected profits from entry under any history.

For example, suppose infringer responses $\{\sigma_2^*(L), \sigma_2^*(\not{L})\} = \left\{Out, \frac{1}{2}\right\}$ are in equilibrium. From $\sigma_2^*(\not{L}) = \frac{1}{2}$, I know his expected profits from entry given $\sigma_2^*(\not{L}) = \frac{1}{2}$, after observing $\{\not{L}\}$ are zero. Therefore, by Lemma 5.1, if the infringer tries to enter with lower probability $\left(\hat{\sigma}_2^*(\not L) < \frac{1}{2}\right)$ upon observing $\{\not L\}$, his expected profit from entry upon observing $\not L$ will become positive, thus forcing $\hat{\sigma}_2^*(\not L) = 1$, contradicting my supposition that $\left(\hat{\sigma}_2^*(\not L) < \frac{1}{2}\right)$. Furthermore, if the infringer tries enter with higher probability $\left(\hat{\sigma}_2^*(\not L) > \frac{1}{2}\right)$ then by Lemma 5.1, his expected profits from entry upon observing $\not L$ will become negative, thus forcing $\hat{\sigma}_2^*(\not L) = 0$, thereby contradicting my supposition that $\hat{\sigma}_2^*(\not L) > \frac{1}{2}$.

By Lemma 6 I know that the only remaining potential alternative equilibria can lie in $\{\{\hat{\sigma}_2^*(L), 1\} \mid \hat{\sigma}_2^*(L) > 0\}$. However I have already shown that $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not L)\} = \{0, 1\}$ will cause the infringer to expect negative profits from entry upon observing $\{\not L\}$. Therefore by Lemma 6 I know that the infringer's expected profits from entry upon observing $\{L\}$ must also be weakly negative when $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not L)\} = \{0, 1\}$. Furthermore, by Lemma 4.2 the infringer's expected profits from entry, conditional on observing $\{L\}$ are continuously decreasing in $\hat{\sigma}_2^*(L)$. Therefore there can be no alternative equilibrium such that $\hat{\sigma}_2^*(L) > 0$. Thus I have shown that if $\{\sigma_2^*(L), \sigma_2^*(\not L)\} = \left\{Out, \frac{1}{2}\right\}$ is an equilibrium, there can be no other equilibrium $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not L)\} \neq \{\sigma_2^*(L), \sigma_2^*(\not L)\}$.

Infringer Best Response Properties

In any equilibrium, the existence of litigation in period 1 will weakly decrease the likelihood of infringement in period 2. This property exists because litigation is positively correlated with the patent holder's private expectation of the likely award in period 2, and thus negatively correlated with the infringer's expected profits from additional infringement. This notion is formalized in Corollary 1 which stems directly from Lemma 6.

Corollary 1 *In any Bayesian equilibrium* $\sigma_2^*(\not L) \geq \sigma_2^*(L)$.

From Corollary 1, I see that litigation must always be a signal that causes deterrence $(\sigma_2^*(L) < \sigma_2^*(\not L))$. Imagine if I assumed instead that $\sigma_2^*(L) > \sigma_2^*(\not L)$. In this case, litigation would be more likely to encourage entry than the absence of litigation, implying that patent holders optimal first period effort would all be monotonically increasing in their type. As a result, this would imply necessarily that $\int [\phi(\theta \mid L)\theta\gamma]\, d\theta > \int [\phi(\theta \mid \not L)\theta\gamma]\, d\theta$.

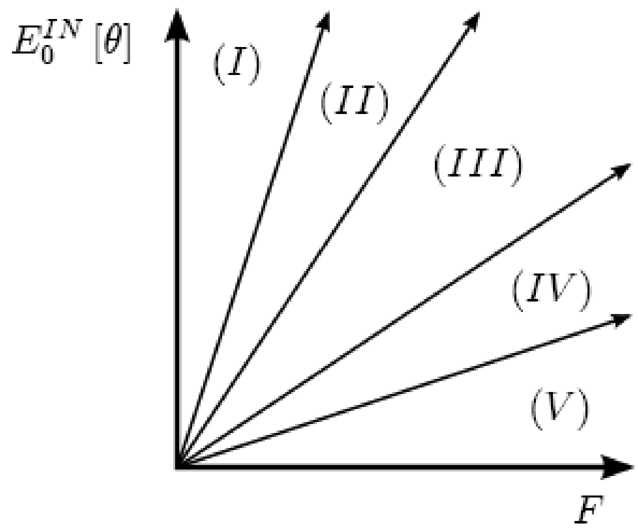

Figure 3: The the resulting class of equilibrium.

This in turn would cause the infringer's expected profit from entry under existing litigation (24) to be less than its expected profits under the absence of litigation (23), which contradicts the hypothetical $\sigma_2^*(L) > \sigma_2^*(\not L)$.

Here I classify equilibria in terms of the entry behavior of potential infringers. This classification system will aid me in the discussion of my results. Thus an infringer's second period behavior within a Class (i) setting is:

(I): always chooses $\{Out\}$.

(II): chooses $\{Out\}$ when observing $\{L\}$ and mixes over $\{In, Out\}$ when observing $\{\not L\}$.

(III): chooses $\{Out\}$ when observing $\{L\}$ and chooses $\{In\}$ when observing $\{\not L\}$.

(IV): mixes over $\{In, Out\}$ when observing $\{L\}$, and chooses $\{In\}$ when observing $\{\not L\}$.

(V): always chooses $\{In\}$.

The Existence of Reputational Effects

In this section, I determine the regions of the parameter space where I find reputation effects. Furthermore, I show that these effects exist whenever the potential infringer's entry behavior is affected by the first period outcome. To this end I define what I refer to as a reputational equilibrium. This will allow me to focus on equilibria where reputations are a factor. After all, any equilibrium where the infringer's period 2 entry action is unaffected by the first period history is trivial in terms of any reputation effects, because the infringer's entry behavior is unaffected by any signal it receives.

Definition 3 *A Perfect Bayesian equilibrium in my model is a* ***reputational*** *equilibrium if*
$\{\sigma_2^*(L), \sigma_2^*(\not{L})\} \notin \{\{Out, Out\}, \{In, In\}\}$.

Here I state the following proposition to show that reputational equilibria exist within a non-empty region of the parameter space.

Proposition 2 *There exists a non-empty region of the parameter space, prior beliefs* $\phi_0(\theta)$, *and type spaces* $\Theta \subset R$, *such that the equilibrium is reputational.*

From the proof of Proposition 2 (in the Appendix), I show that the equilibrium will depend on the size of ex-ante expected awards $E_0^{IN}[\theta] = \int_{\theta\in\Theta} \theta\phi_0(\theta)\, d\theta$ relative to the size exogenous licensing fee F. Consider Figure 3. Here I demonstrate that the existence of non trivial equilibria (Regions II,III,IV) relies on the ex-ante expected award being neither too large in comparison to the licensing fee (Region I), nor too small (Region V).

The Properties of Optimal Effort

In this section I uncover the relationship between the incentives for reputational effects and a patent holder's first period effort. To this end, I use Propositions 4 and 5 to show the thresholds in the space of private information where one effect (deterrence or luring) begins to dominate the other.

Proposition 3 *In any reputational equilibrium, there exists* $\overline{\theta} \in R$ *such that first period effort is*
(3.1) *increasing in* θ, *if and only if* $\theta < \overline{\theta}$
(3.2) *decreasing in* θ, *if and only if* $\theta > \overline{\theta}$.

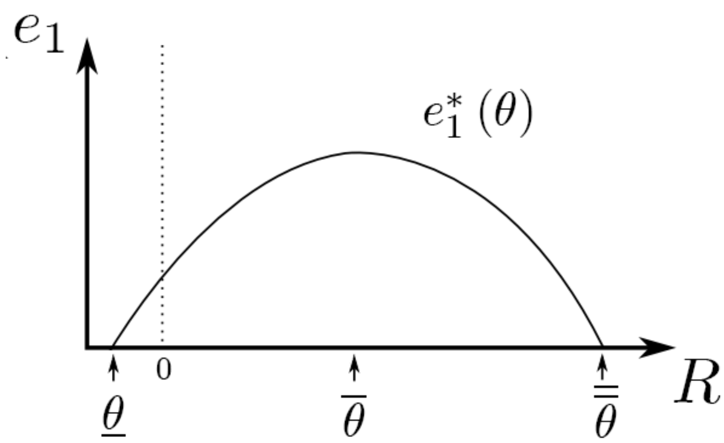

Figure 4: Optimal patent holder effort in a reputational equilibrium.

Proposition 4 *In any reputational equilibrium, there exists* $\overline{\overline{\theta}} \in R$ *and* $\underline{\theta} \in R$ *such that*

(4.1) $e_1^*(\theta) = 0$, *if* $\theta < \underline{\theta}$

(4.2) $e_1^*(\theta) = 0$, *if* $\theta > \overline{\overline{\theta}}$.

Propositions 3 and 4 give me thresholds in the patent holder type space. These thresholds ($\underline{\theta}, \overline{\theta}$, and $\overline{\overline{\theta}}$) are used in Figure 4 to compare each patent holder type's first period effort. I see here that first period effort is non-monotonic in type. This property occurs because, for all types θ greater than $\overline{\theta}$, the benefits of deterrence are diminishing, relative to the benefits of luring, as I increase θ.

Note that for all patent holder types $\theta \in \left[\overline{\theta}, \overline{\overline{\theta}}\right]$, the patent holder is still choosing positive effort levels in period 1. Therefore, if a patent holder of this type succeeds in initiating litigation in period 1, they will still go to trial. Furthermore, all patent holder types $\theta \geq \overline{\overline{\theta}}$ are expending zero effort in period 1, even though they have the highest award expectations, because they do not want to deter entry under in period 2 under any circumstance. This behavior may be characteristic of what Herald (2008) refers to as a Patent Troll: A patent holder that relies on litigation to extract revenue from potential licensors, who obscures their own property rights to encourage infringers.[17]

[17] An endeavor to find observable statistics that identify potential patent trolls post litigation merits consideration.

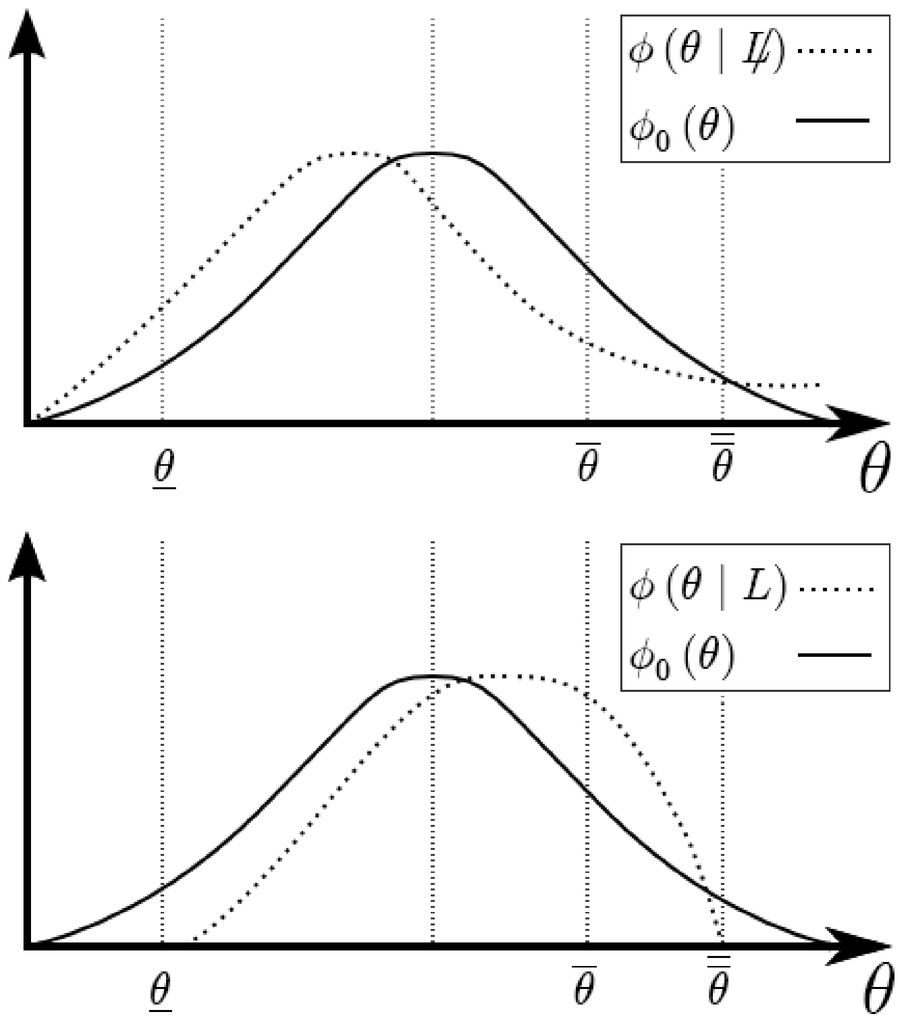

Figure 5: Changing beliefs in an example setting that results in a reputational equilibrium.

Corollary 2 *The probability that litigation occurs in period 1, is weakly increasing in* θ*, if and only if* $\theta < \overline{\theta}$*.*

Corollary 2 is a direct result of Proposition 2 and corollary 1. When a patent holder is benefiting from the deterrence effect, increasing its expected award makes it more likely that he will create litigation in period 1. On the other hand, when a patent holder is benefitting from the luring effect, increasing its expected award makes it less likely that he will create litigation in period 1. This non-monotonic property is illustrated by the optimal effort depicted in Figure 4. Note that some patent holders are so interested in luring the potential infringer, that the probability that they create first period litigation is zero (types $\theta > \overline{\overline{\theta}}$).

Consider figure 5. This figure shows the change in beliefs, for each first period history, with respect to the thresholds defined in Propositions 3 and 4. I show in Figure 5 that settings exist where the absence of litigation from period 1 causes the infringer to enter, even when it assigns a higher probability density to some patent holder types with the highest award expectations.[18] At the same time, the existence of litigation will deter entry, even though the probability of patent holder having the highest award expectation goes to zero, because all patent holder types greater than $\overline{\overline{\theta}}$ will never create litigation in period 1.

Corollary 3 *The magnitude of reputation's effect on patent holder effort is*
(C3.1) *increasing as* $\sigma_2^*(\not L) \to 1$
(C3.2) *decreasing as* $\sigma_2^*(L) \to 1$

In any trivial equilibrium, the infringer's entry action is unconditional; therefore, patent holders can not benefit from either deterrence or luring effects, because the infringer's action is unaffected by the signal they observe. Thus, I can analyze how the conditionality of the infringer's entry decision affects the patent holder's incentives to deceive, by observing the relative changes in patent holder effort as the equilibrium infringer responses move from $\{\sigma_2^*(L), \sigma_2^*(\not L)\} = \{Out, In\}$ (the dotted line in Figure 6)

[18] This effect on beliefs is analogous to a poker game, where a player passes on the opportunity to raise the pot. Opposing players might assign a higher probability to them having a weak hand, and at the same time assign higher probability to having a really strong hand.

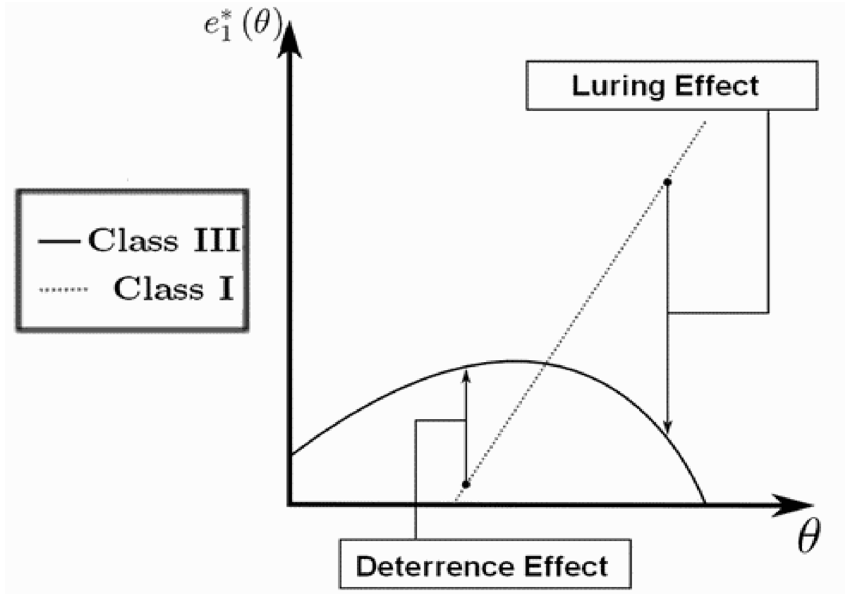

Figure 6: Reputation's effect on patent holder effort.

to $\{\sigma_2^*(L), \sigma_2^*(\not{L})\} = \{Out, Out\}$ (the solid line in Figure 6), which could be caused by a decrease in the licensing fee. The relative changes in patent holder effort are a result of Lemma 3, which shows the change in patent holder effort across types, for an increase in $\sigma_2^*(\not{L})$.

As illustrated in Figure 6, when moving from a conditional entry equilibrium (Class III) to an unconditional entry equilibrium (Class I), the optimal effort in period 1 becomes steadily increasing in type, because no type wants to expend more or less effort for the sake of any reputation effect. Furthermore, I observe that for all types greater than $\overline{\theta}$, patent holders are spending less effort when infringer's entry decision is conditional, than if it were unconditional. These types are doing so because the benefits from luring are outweighing the benefits from deterrence. At the same time, I observe that below $\overline{\theta}$ in the type space, patent holders are spending more effort. These types are doing so because the benefits from deterrence are outweighing the benefits from luring.

1.4 Discussion

Implications of Award Distributions The existence reputational equilibria depends on both the expected award from litigation $E_0^{IN}[\theta] = \int_{\theta \in \Theta} \theta \phi_0(\theta)\, d\theta$ and the variance of the expected award. Distributions, with relatively low means ($E_0^{IN}[\theta]$) and long and thin upper tails create the ideal setting for predatory behavior because most patent holder types will try and create deterrence. In these settings, the absence of litigation in period 1 will be very enticing to a potential infringer. Thus, predatory patent holders will be able to use the luring effect to good advantage, by delaying litigation. On the other hand, consider a distribution with two symmetrically distributed types: a patent holder type with extremely large award expectations and those who expect zero awards from litigation. In this setting, patent holders with extremely large award expectations are effectively unable to use the luring effect, because the there is no one else for the zero expectation patent holders to pool with (to create deterrence by imitation). For this reason, when there are only two types, I know that the first period effort of the low expectation type must be less than the first period effort of the high expectation type. This type of setting will usually result in equilibria with minimal opportunities for deterrence and luring (either Class II or Class IV depending on the prior density of the two types).

I have shown that for any nontrivial equilibrium that there exists a type $\overline{\theta} \in R$ such that for all $\theta > \overline{\theta}$, $\frac{\partial e_1^*(\theta)}{\partial \theta} < 0$. However to show that $\overline{\theta} \in \Theta$, requires more specific assumptions about the distribution. For example, under certain parameter settings it is possible that the patent holder with the largest award expectation is threatening to a potential infringer, but doesn't want to encourage entry. In these cases, I see reputational behavior analogous to the findings in Kreps and Wilson (1981), where every patent holder has some incentive to create deterrence.

Policy

A policy maker might wish to control the impact of reputation in patent litigation. Furthermore, I have shown how these incentives are linked to the distributions of awards.

Therefore, since policy makers can affect the distribution through the average size of potential awards and the requirements for recovering those awards, I conclude they have some control over reputational behavior.

For example, allowing for awards in excess of 3 times damages in special circumstances might stretch out the upper tail of an existing distribution of expected awards. From Proposition 3 I know that this will weakly increase the likelihood of predatory behavior (Luring). On the other hand, relaxing the requirements for receiving modest awards may shift the bottom part of the distribution towards the middle. From corollary 3 I see that this will weakly decrease the amount of anti-competitive behavior (Using deterrence when perfect information about the patent holder's property rights would lead to entry).

In general, vagueness about a patent holder's property rights will broaden the distribution of possible expected awards. In this sense, the reputational distortions present in my model may represent a cost associated with a lack of clarity in patent law. At this point I do not conjecture about the actual costs to society of any reputational behavior. These questions lay beyond the scope of my model in its current form. However, with the proper additions to my model I believe I can specifically address policy issues. This may present an interesting avenue for research in both theoretical and empirical analysis.

Endogenous Licensing Fee

To this point I have assumed that the licensing fee is exogenous for simplicity. If I were to allow the licensing fee to be endogenous, I would expect it to be a function of the observed history and the patent holder's private information.

I show that if I apply the following common sense restrictions on the endogenous nature of the licensing fee, I do not change the general intuition of my original model.[19]

Assume the licensing fee $F(\theta, h_1)$ is a twice continuously differentiable function of θ

[19] Under the following assumptions, Lemmas 1, 2, and 3 are satisfied for an endogenous licensing fee. These are sufficient conditions to prove Propositions 1 and 2.

and h_1 such that:

$$\frac{\partial F(\theta, h_1)}{\partial \theta} > 0, \tag{29}$$

$$F(\theta, L) > F(\theta, \not{L}) \tag{30}$$

$$F(\theta, L) > \int_{\theta \in \Theta} (\theta\gamma)\, \phi_0(\theta)\, d\theta \tag{31}$$

$$F(\theta, \not{L}) < \int_{\theta \in \Theta} p\left(e_2^*(\theta, In)\right)(\theta\gamma)\, \phi_0(\theta)\, d\theta \tag{32}$$

The intuition of these restrictions is as follows. In (29) I require that any licensing fee secured by the PH is increasing in their expected award. In (30) the licensing fee gained under the threat of current litigation must be greater than the licensing fee obtained without ongoing litigation. In (31) the licensing fee agreed upon in the presence of litigation is greater than the expected litigation award for the additional act of infringement given initial beliefs. In (32) the licensing fee agreed upon in the absence of litigation is less than the expected litigation award for the additional act of infringement given initial beliefs.

1.5 Conclusions

I have developed a two period model where litigation timing affects the endogenous entry behavior of potential infringers. My model adds to the existing literature by considering the timing of litigation in the enforcement of patents and addressing the delay of information revelation inherent in the long litigation process. Furthermore, I have uncovered theoretical motives that can explain why patent holders with really high expectations of litigation awards might strategically delay litigation. This finding could potentially be used by firms found guilty of infringement to argue for a reduction in awards whenever there is a noticeable delay between an initial act of infringement and the beginning of litigation. Also, I provide a mechanism for policy makers to affect reputational behavior.

With regards to signaling games in general, my model presents an environment where there are two opposing reputational effects. Patent holders with weak property rights might want to appear strong. At the same time, patent holders with strong property rights might want to appear weak. Similar phenomenon may exist in other signaling games that have the potential for conflicting reputational effects. In this paper the high types actually benefit from their type being underestimated. One could imagine many other situations like this. For example, a ratchet effect model where perhaps intermediate types want to show off to keep their jobs but high types want to avoid showing off initially to avoid a ratcheting up of the employer's demands.

CHAPTER 2

WELFARE IMPLICATIONS OF ONLINE PUBLISHING

2.1 Introduction

Advances in information technology have greatly increased the variety of goods available to consumers. While the consequences for consumer demand and welfare have been the subject of much research, the effects on producers, and potential entrants in particular has not. We study the effect of a new book publishing technology on entry in a segment of the book industry for which data on both sales and costs was available. Although this segment is relatively small, the ease with which data can be collected and simple pricing terms makes it attractive for study and our qualitative conclusions extend to other more economically significant examples of user generated content.

When an author is unable to find a publisher willing to take on the financial risks of printing and distributing their book he or she can pay a vanity publisher to have the book printed. The term vanity publishing reflects the fact that these books lose money and their authors are purchasing the pride of seeing their work in print. Huberman, Romero and Wu show that YouTube video uploaders derive some non-pecuniary benefit from having their videos viewed by others by demonstrating a positive relationship between increasing viewership and future productivity of YouTube uploaders. Our analysis seeks to estimate the parameters governing the distribution of vanity in a population of authors considering vanity publishing.

Up until ten years ago, all vanity publishers employed offset printing to publish their books.[20] With offset printing, the initial portion of almost every print run is unusable because the ink is not evenly distributed on the printing plates until the press has been

[20] Offset printing is a technology where the inked image is transferred (or "offset") from a plate to a rubber blanket and then stamped on the printing surface. There are substantial fixed costs because new printing plates must be created and loaded into the press for each book published.

running for a while. Consequently, the minimum print run with traditional offset technology is hundreds of copies. In contrast, publishing on-demand (POD) allows publishers to print commercially competitive books a single copy at a time.[21]

Much prior research has considered the effects of technological change on creative-content industries. With data on a large number of similar products, all of which have been subject to the same technological change, researchers can identify the effect of the change on existing products. For example, Blackburn (2004) estimates the impact of file sharing on the music industry and Brynjolfsson, Hu and Smith (2003), henceforth BHS, estimate the effect on consumer welfare due to emerging Internet sales technology in the book industry. In contrast to BHS, who focus on consumer welfare gains from increased product variety in the book market, our research estimates the welfare gain of producers from the emergence of POD publishing technology. We use results reported by BHS to impute books sales from Amazon sales ranks for a subset of authors who publish using POD to parameterize an expected utility function for a population of authors, and model their decision to publish.

Our innovation is to apply tools from the patent valuation literature to entry in a creative-content industry. We do not observe the entire value to the author of seeing their work published, just as the value of a patent to its holder is not directly observable. Prior to publication, each author faces uncertainty over how many copies their book will sell. This uncertainty is resolved over time in the same way as a patent holder's uncertainty over the value of the patent. Because POD publishers compete with one another for paying customers, their pricing is as transparent just as patent office fees are. Identification of the parameters governing the distribution of patent valuations is achieved in Pakes (1986) by matching the simulated patent renewal rates to the rates observed. We will identify the parameters governing the joint distribution of author's valuations and expected sales by matching the simulated distribution of sales to the observed sales distribution.

Authors would set lower prices than on demand publishers do because authors of

[21] Publish on demand employs digital printing. Book covering and binding can be achieved inline if the digital press has a binding module, although small batch stand alone book binding equipment is available.

obscure content typically consume a nontrivial fraction of total sales (an equivalent result can be obtained if authors derive some non-pecuniary marginal benefit from each sale). This conclusion extends to other markets where user generated content is delivered via platforms such as YouTube or MySpace. As is typical of "free" content, consumers pay for content by being subjected to advertising, and the amount of advertising the platform chooses is higher than the content's creator would prefer.

The rest of the paper is organized as follows. In Section 2 we discuss the publishing industry in the context of recent research. We model the author's utility maximization problem in Section 3. In Section 4 we describe the simulation methods we use to estimate parameters governing the distribution of expected sales and author utility. In Section 5 we describe our data and the method we use to impute sales from sales rankings. We present our results in Section 6, and Section 7 concludes the paper.

Publishing Industry

Since Gutenberg, there have been significant economies of scale in printing. This, combined with the significant costs of editing and typesetting, has made it nearly impossible for publishers to turn a profit on low volume books. Limited shelf space constrains bookstores to offer only a fraction of the universe of available titles.[22] Bookstore owners rely on a variety of signals of book quality when deciding which books to stock (e.g. book reviews, sales of previous works by the same author, endorsements from well-known individuals on the back cover, return policy, etc.). In this environment, a publisher's decision to invest in the production and distribution of a book is a powerful signal of quality because bookstores may return unsold copies to the publisher for a refund in the first year after publication. While the refund may be partial, allowing returns helps to align the publisher's incentives with bookstores'.[23]

Most POD books are available for sale exclusively online or from the author. Without the credibility a non-subsidy press conveys, the barriers to bookstore distribution

[22] A large bookstore might stock 20-40 thousand unique titles whereas there are ~2.3 million books in print. Brynsdolf, Hu and Smith.

[23] If a publisher refused to accept returns it would always have a strong incentive to sell more copies to bookstores and so would be unable to send a creditable signal of expected sales.

are substantial. Anecdotal evidence suggests that most bookstore owners will not even consider stocking a self-published or vanity press title unless the author convinces them otherwise in person.[24] Only 0.25% of the titles in our dataset ever achieved substantial retail distribution, so our decision to ignore brick-and-mortar retail sales should be justified for the remaining 99.75%.

We focus on a population of authors who are unable to have their current manuscript picked up by a legitimate publisher. There are many reasons why potential authors are unable to find a legitimate publisher. New authors may lack the skill, marketability, or reputation to justify publication. New authors may be unable to find the appropriate channels for distribution, or afford the search costs for marketing their work to potential publishers. Also, some authors may just be unlucky, in the sense that even though there work might be highly marketable, asymmetric information in the review process, and a failure to signal their true talent causes them to be passed over. Even so, given the small number of successful commercial authors relative to the number of potential authors as a whole, one could imagine that the distributions that govern expectations and vanity of the subset of vanity authors is still representative of authors as a whole.

The Long Tail

The internet offers consumers both a larger variety of products and tools that lower the cost of search. Whereas niche products are typically unavailable at most retail outlets, on the internet they sometimes make up a large fraction of sales, a phenomenon Anderson (2004) coined the Long Tail. The relative importance of blockbusters and obscure products, (i.e. the head and tail of the product distribution) has been a subject of considerable debate between those who believe that "the tail is likely to be extremely flat and populated by titles that are mostly a diversion for consumers whose appetite for

[24] See e.g. Writer Beware, Science Fiction and Fantasy Writers of America (http://www.sfwa.org/Beware/), or the website of Johnathan Clifford who claims to have coined the term “Vanity Publisher” (http://www.vanitypublishing.info/). Of the more than 17 thousand titles published by iUniverse by 2004 only 20 made it onto Barnes & Nobles shelves, despite the 25% ownership stake Barnes and Nobles held in iUniverse. XLibris authors did little better; only 20 of the 8 thousand titles published by 2004 were picked up by traditional publishers for retail distribution. Gail Feldman “Got a Book in You?”, New York Times, March 1, 2004.

true blockbusters continues to grow"[25] and those who believe consumer demand "will continue to shift from a few best-selling products to niche products"[26]. BHS estimate that 13 million copies of books not available at typical brick and mortar retailers are sold each year, whereas POD account for some ten or twenty million books over the last decade, and J.K. Rowling's Harry Potter series has sold several hundred million copies worldwide, all since 1997.[27] In the book industry blockbusters continue to dominate.

BHS examines the impact of increased product variety made available through electronic markets. Specifically, they calculate the increase in consumer surplus resulting from books sold at online bookstores, which are not available at traditional "brick and mortar" bookstores. They dub these newly available books "obscure titles." POD books clearly fit their definition of obscure books as traditional bookstores almost never stock titles from subsidy publishers. Their strategy is to assume that books with Amazon.com sales ranks greater than some cutoff value are only available online. Then, they estimate the relationship of sales rank to actual weekly sales and integrate over all books with sales ranks above the cutoff. In comparison to the effect of newly available but previously created content, content specifically created in response to the increase in variety is small.

What is Vanity?

We define vanity as the portion of utility gained by an author from the publication of his or her manuscript which is independent of any profits. According to the NEA 7% of American adults engage in creative writing at least once per year, and approximately one seventh have their works published.[28] Most artists want an audience for their work; similarly, most people want attention from others. As with gambling or playing sports or accumulating collectibles, in the arts there is a spectrum of talent and potential income: most amateurs pay to enjoy their hobby, and some individuals are able to make a small

[25] Elberse (2008).

[26] Brynjolfsson, Hu and Simester (2007).

[27] Brynjolfsson, Hu and Smith report obscure books sales of $578 million per year and an average price of $42.18, implying sales of some 13 million books.

[28] National Endowment for the Arts, 2002 Survey of Public Participation in the Arts.

income, but only the most talented are able to make a living at it.

None of our results depend on the precise nature of the utility gain, so long as the author's welfare depends only on money income and the binary decision, publish or don't publish. We rely on variation across potential authors in income, enjoyment of writing and egotism, to generate a distribution of utility from publication which we refer to as vanity. We favor the traditional explanation implied by the term "vanity press"; authors pursue the pride resulting from seeing their work in print. Of course, many individuals who do not think of themselves as authors write books. Having written a published book on a relevant topic to one's profession is an important credential that would likely appear on a resume. Amateurs who hope to become professional authors might hope that the notice (or at least feedback) they receive for their current book may increase the chance of success in the future. However, given the extreme improbability of commercial success, the chief source of utility is being able to think of oneself as a "published author." An alternative theory that could potentially explain our observations about the non-pecuniary benefits of using POD could be that authors are investing in human capital. To counter this claim would require a thorough analysis of the sales performance of these authors over their lifetime. However, there are a number of authors who have published more than one book with POD, and there does not appear to be a significant increase in sales from an author's first title to later titles.

If authors care about how many copies of their books are sold independently of the effect of sales on profits then our results will be biased. xLibris does sell marketing services for the books it publishes.[29] Although anecdotal evidence suggests that bookstores are entirely unresponsive to this marketing, authors might be purchasing marketing because they gain additional vanity based utility from being a "real" author whose books are carried by brick and mortar retailers. As would be expected if their marketing services were as ineffective as we believe, POD firms do not publicly release data pertaining to the effectiveness of their marketing services.

[29] For example, an advertisement in ForeWord, a magazine containing reviews of new books published by independent presses, costs $399.

2.2 The Model

In this section we develop a model of author choice that incorporates the possible decisions facing authors at the time our data was collected. The goal is to establish a theoretical framework that simulates a potential author's behavior given their individual preferences and expectations. The results of this section will give insight into the distributions of potential author preferences and expectations, by matching the observed characteristics of authors choosing to publish with Print on Demand, with the characteristics of simulated authors choosing to publish with print on demand.

Players

We envision a population of authors, each of whom has written a single manuscript. Individual potential authors are characterized by their demand functions and vanities. Let an author $i's$ random demand function be $x\left(s_i, \varepsilon_i\right)$, where s_i is the author's private information and ε_i is a random component of realized sales that remains unobserved until the author chooses a publication method. We assume $s \perp \varepsilon$. Let an author's vanity be v_i, which is the value they receive upon publication of their book. Therefore, author types are drawn from the corresponding distributions $f(s)$, $g(v)$ defined over respective domains S and V. Therefore an author type t is a member of the set $T = \{S \times V\}$. We denote the joint distribution of types $\tau\left(s, v\right) = f(s)g(v)$.

Actions

An authors action profile a, is defined by their publishing decision such that $a \in A = \{None, POD, \{TRD_Q\}\}$ is a choice between respective actions: don't publish, Print on Demand, and traditional vanity press, with various quantity choices such that $\{TRD_Q\} = \{200, 500, (1000, \infty)\}$.

Utilities

An author $i's$ utility $\mu(a \mid s_i, v_i)$ depends on their type, action, realization of sales, royalty rates, and costs.[30] Therefore we denote the author's utility for each action respectively as:

$$
\begin{aligned}
\mu(none \mid s_i, v_i) &= 0 \\
\mu(POD \mid s_i, v_i) &= v_i + R_{POD}x(s_i, \varepsilon_i) - C_{POD} \\
\mu\left(TRD_{Q_j} \mid s_i, v_i\right) &= \left\{ \begin{array}{ll} v_i + R_{Tra}x(s_i, \varepsilon_i) - C_{Tra}(Q_j), & \text{if } x(s_i, \varepsilon_i) < Q_j \\ v_i + R_{Tra}Q_j - C_{Tra}(Q_j), & \text{if } x(s_i, \varepsilon_i) \geq Q_j \end{array} \right\}
\end{aligned}
$$

C_j and R_j are the cost and author's royalty rate under technology j, and $\{Q_j, C_{Tra}(Q_j)\}$ are the menu of available traditional print-run sizes and costs. Note that if the author chooses traditional vanity publishing, the author must choose the print-run before publishing and so the number of copies printed is a capacity constraint and is therefore affected by realization of demand.

2.3 Simulation

Simulated Authors

We take create a simulated potential author by taking separate random draws s_i and v_i from the distributions $f(s)$ and $g(v)$. We repeat this process n times to form a population $P = \{s_i, v_i\}_{i=1}^{n}$ of n potential authors. If i author chooses to publish a book, we then add noise by taking a random draw ε_i from the distribution $h(\varepsilon)$ such that simulated sales for an author i becomes $x(s_i, \varepsilon_i) = s_i\varepsilon_i$. In order to estimate our model we need to make assumptions about functional forms of the distributions governing expected sales, vanity, and uncertainty.

Assumption 1 $s \sim exp(\alpha + \beta * log(Uniform[0, 1]))$ *Pareto*

Assumption 2 $v \sim exp(N(\mu, \sigma_v^2))$ *lognormal*

[30] xLibris's pays 10% on all sales at Amazon, so we assume a constant 10% royalty rate.

Assumption 3 $\varepsilon \sim \exp(N(\frac{-\sigma_{\varepsilon}^2}{2}, \sigma_{\varepsilon}^2))$ *lognormal.*

We seek to estimate the five parameters $\theta = (\mu_v, \sigma_v^2, \alpha, \beta, \sigma_\varepsilon^2)$: the mean and variance of the author's utility from being published, μ and σ^2 respectively; the parameters governing the mean and variance of the distribution of author's expected sales, α and β respectively; and the variance of the ratio of realized sales to author's expected sales, σ_ε^2.

The simulation requires that authors choose the action that maximizes their expected profits. Given the functional form of the above distributions as determined by the set of unknown parameters θ we denote the expected profits for an author type $t_i \in (S \times V)$ as

$$
\begin{aligned}
\mu^E\left(none \mid t_i, \theta\right) &= 0 \\
\mu^E\left(POD \mid t_i, \theta\right) &= v_i + s_i R_{POD} - C_{POD} \\
\mu^E\left(TRD_{Q_j} \mid t_i, \theta\right) &= v_i + \int_0^{\frac{Q_j}{s_i}} s_i \varepsilon R_{Tra} h\left(\varepsilon \mid \sigma_\varepsilon^2\right) d\varepsilon + \int_{\frac{Q_j}{s_i}}^{\infty} Q_j R_{Tra} h\left(\varepsilon \mid \sigma_\varepsilon^2\right) d\varepsilon - C_{Tra}\left(Q_j\right)
\end{aligned}
$$

Because POD ensures that no unsold copies are produced, expected profit under this printing technology is unaffected by the variance of Noise(σ_ε^2). We assume risk neutrality. If authors are risk averse, then expected utility from a POD book will be decreasing in the variance of noise. However, assuming risk aversion, expected utility from traditional publishing would decrease even faster.

Figure 1(Following Page) illustrates an author's optimal choice as regions in the type space. Notice that the region where types are choosing not to publish depends on both their vanity and expected sales. Author types with relatively small vanities may choose not to publish even thought their expected sales are greater than other types with lower expected sales and higher vanities. As a consequence, the distribution of observed POD sales will depend also on the distributions of vanity. This property will help us identify the parameters of the distributions of vanity and all authors expected sales that result in a distribution of POD sales similar to the distribution of our data. Figure 2(Following

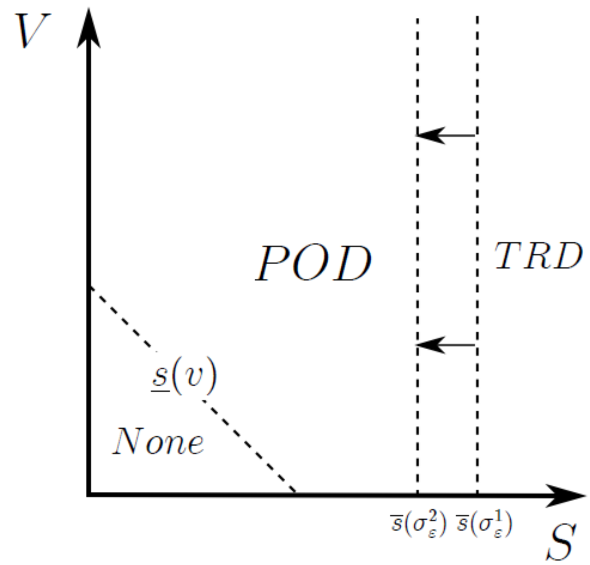

Figure 7: Showing a shift in the boundaries of potential author choices for a decrease in σ_ε.

Page) illustrates level sets for expected utility in the type space of author characteristics for two publishing decisions: POD and a traditional vanity printing of 200. The curvature of the level set for a traditional vanity publishing choice reflects the tension between diminishing average costs and a capacity constraint. In contrast, the level set for POD is a straight line, reflecting the notion that since noise has no impact on expected POD sales, vanity and expected sales are perfectly substitutable(in terms of expected utility).

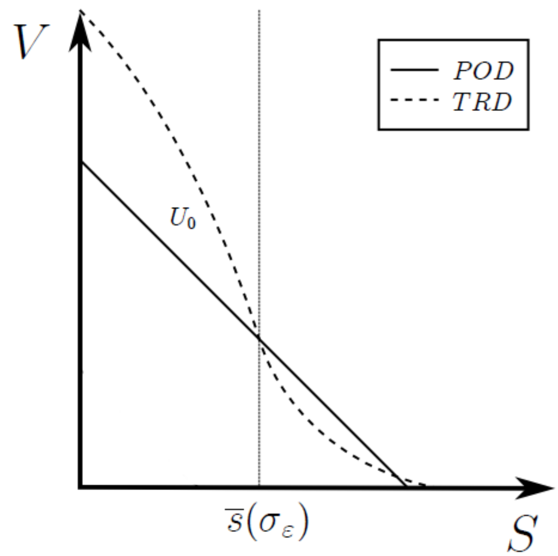

Figure 8: Level sets for utility (U_0) in the typespace of author characteristics for two publishing technologies POD and Traditional ($Q = 200$).

2.4 Estimation

We identify our parameters by functional form. Using our understanding of an author's decision making process, we can define the space of author types that will choose POD as a function of an individual author's vanity and expected sales.

Let $T_{POD}(\theta) = \{t' \in T \mid a^*(t' \mid \theta) = POD\}$. Therefore we define the distribution of sales for authors choosing types POD as

$$z(\varepsilon s \mid \theta) = \tau(s', v' \mid s', v' \in T_{POD}(\theta))\, h(\varepsilon).$$

Our simulation exercise allows us to estimate the parameters θ by finding values $\hat{\theta}$ such that the conditional distribution of sales of simulated authors $z(\varepsilon s \mid \widehat{\theta})$ resembles the the distribution of observed POD sales in the data.[31] We choose estimated parameter values

[31] This procedure is analogous to a GMM estimation in which not all the moment restrictions can bind with equality; every sales level at which the simulated and observed CDF must be equal puts another restriction on the parameters. Because not all these restrictions can be met simultaneously (that is, no parameter values will make the simulated and observed distributions exactly equal at each sales level), whatever metric we use to evaluate how "close" the two distributions are will effectively be some function that assigns weights to the restriction that the CDF of the simulated and observed sales distributions are equal at some sales levels.

minimize the mean squared difference in CDFs at each level of observed sales in the data.[32] Our results did not change when we used the mean, absolute difference in CDFs. Therefore we minimize the sample analog of the mean squared difference in CDF.[33]

Let $Z_{Data}\left(Sales\right)$ be the cumulative distribution of observed sales in the data. Let $Z(\varepsilon s \mid \widehat{\theta})$ be the cumulative distribution of simulated sales for all simulated types $t \in P$ such that $a^{*}(t \mid \widehat{\theta}) = POD$. Therefore, our objective function becomes

$$\int_{0}^{\infty} \left[Z_{Data}\left(Sales\right) - Z(Sales \mid \widehat{\theta})\right]^{2} d(Sales) \approx$$

$$\sum_{i=1}^{N}\left(\frac{\#\left(ObsSales \leq ObsSales_{i}\right)}{N} - \frac{\#\left(\text{Sim}Sales \leq ObsSales_{i}\right)}{N}\right)^{2}.$$

which we minimize w.r.t the parameters $\theta{=}(\mu_{v}, \sigma_{v}^{2}, \alpha, \beta, \sigma_{\varepsilon}^{2})$.

For every simulated population of potential authors we can use our choice model to determine the sub-population that will choose to publish using POD. Using the population of POD authors, we can then simulate the distribution of realized sales. Identification occurs because each parameter to be estimated has a different impact on this simulated distribution of realized sales at an estimated parameter value. In other words, we cannot replicate the impact of a change of one variable on the cumulative densities at all levels of sales, by holding that variable fixed and instead using some variation of the other parameters. The intuition for identification is presented at length in Appendix B.

Once we have estimated parameters values for a simulation run, we can compute the probability that an author will choose POD publishing (it is the fraction of the n simulated authors who choose POD publishing under the estimated parameter values). One over this probability is the number of manuscripts per POD published book, so dividing N, the observed number of POD books, by the estimated probability that an

[32]Note that the distribution of expected sales in our author population is not the distribution of expected sales for POD authors. POD authors are a subset of this overall population.

[33]To find the minimum of our objective function, we rely on the MATLAB algorithm fminsearch which uses the Nelder-Mead simplex search algorithm. Since proving the uniqueness of the solution analytically is impractical, we instead vary the programs initial guess(using 1000 randomized settings for our first simulation run) to make sure that it converges to the same minimum.

author will choose POD publishing yields the number of estimated number of potential entrants.[34]

We report errors that account for both simulation error inherent in the simulation process, and sampling error. These errors are computed by sampling with replacement from our data of authors to create B bootstrap samples of size w. For each bootstrap sample b, we repeatedly taking new random draws from the underlying distributions, and re-estimate our parameters for T simulation runs. Thus for each sample b we find the estimated parameter values $\left(\hat{\theta}_1^b, ..., \hat{\theta}_T^b\right)$, determine the mean estimate $\hat{\theta}^{*b} = \sum_{t=1}^{T} \hat{\theta}_1^b$, and then calculate the simulation error as:

$$s^2_{\hat{\theta}_b,\mathrm{Sim}} = \frac{1}{T}\sum_{t=1}^{T}\left(\hat{\theta}_t^b - \hat{\theta}^{*b}\right)^2.$$

We then account for sampling error by incorporating the average parameter estimate for each bootstrap replication. Thus the measure of sampling error becomes

$$s^2_{\hat{\theta},Boot} = \frac{1}{1-B}\sum_{b=1}^{B}\left(\hat{\theta}^b - \overline{\hat{\theta}^{*b}}\right)^2.$$

where $\overline{\hat{\theta}^*} = \sum_{b=1}^{B} \frac{1}{B}\hat{\theta}^{*b}$.

We report two sets of standard errors: First, we calculate standard errors assuming that our imputed data parameters are precisely identified. Then as a robustness test, we calculate standard errors that take into account the standard errors reported in BHS. These "two-stage" standard errors allow us to measure the robustness of our results to variation in the assumed parameter values. Therefore, we take separate draws, s_1 and s_2, from two normal distributions with standard deviations equal the standard errors reported in BHS . Then, we proceed as if the true values of β_1 and β_2 differed from BHS's estimates by s_1 and s_2 standard deviations respectively. We take n new draws, carry out both the estimation procedures described above, the welfare calculations described below and store the results. This procedure is repeated T times for each of the B bootstrap samples. The resultant distribution is the basis for our standard errors reported in Tables

[34]A more detailed explanation of this calculation and an example can be found in Appendix B.

A.1-A.4.

Our results depend on the assumed functional form of vanity. To illustrate this we present results for alternate assumptions in Tables C, D, and E in the Appendix. In particular, the estimated population of potential authors is particularly sensitive to our assumed form of vanity. Allowing for negative values of vanity, such as the case with a normal distribution decreases the number of stimulants that end up publishing, which it. In a similar manner, allowing for a vanity distribution with fixed mass points at zero will also increase the estimates of author population relative to our initial assumptions. Furthermore, we find that increasing the mass at zero will further increase estimates of potential author populations, although with relatively little impact on the estimates of other parameters. However, since our welfare estimates are derived from the number of current POD authors that would not have published but for the existence of pod technology, and since the conditional distributions of simulated authors who choose to publish with POD varies less than estimates of the overall population, our welfare calculations are less sensitive to our assumptions of functional form.

2.5 Data

On November 11, 2004 we collected data from Amazon.com on books published by the two largest POD publishers: xLibris and iUniverse. Our data were collected at a time when these two leading publishers dominated the POD market. Although several other POD publishers existed at the time none has more than a few hundred books on Amazon. The subsequent entry of more niche focused and substantial competitors eventually led iUniverse and xLibris to announce a merger in January 2009.[35] Their merger should relieve competitive pressure on the firms, because both targeted the same segment: amateurs who wanted to think of themselves as "professional" authors. We observe all titles available on Amazon.com published by either firm.

[35] Booksurge was a small POD until its acquisition by Amazon in 2005. Lulu.com and CreateSpace (another Amazon company) take the POD concept to its logical conclusion charging no upfront fees and providing distribution and on demand production of CDs, DVDs and electronic content. The author need only purchase a single proof copy before their written, audio or video content becomes available for sale on Amazon. Because the process is entirely automated, Lulu.com and CreateSpace have almost zero fixed cost per title and makes a profit even if the only copy ever sold is the author's proof.

In the text, we present xLibris data and results. Since xLibris and iUniverse offer virtually identical services at the same costs and posses similar market shares(xLibris 43% and iUniverse 37%) over the time period we examine, we use the results from our analysis of the xLibris to calculate the welfare effects for the entire market. In other words, we assume that xLibris and iUniverse are perfect substitutes for authors choosing POD publishing, and scale our welfare results according to the firms' respective market shares.

According to Chevalier and Goolsbee, "Amazon claims that for books in the top 10,000 ranks, the rankings are based on the last 24-hours and updated hourly. For books ranked 10,001-100,000, the ranks are updated once per day. For books ranked greater than 100,000, the sales ranks are updated once per month (Amazon, 2000). Many hundreds of thousands of books, however, have a rank but almost certainly have less than one sale per month. Italie (2001) claims that for these rarely purchased books, Amazon bases the rank on the total sales since Amazon's inception." None of the books in our sample have a rank under 10,000 the vast majority have ranks greater than 100,000.

We construct a measure of sales since Amazon's inception using sales rank and BHS's results.[36] Using data obtained from a small publisher on 321 obscure books, they fit their data to a log-linear distribution:

$$Log(Quantity) = \beta_1 + \beta_2 * log(rank) + \varepsilon$$

and found

$$\beta_1 = 10.5(.156) \qquad \beta_2 = -0.87(0.017).$$

We use these estimates to impute actual sales from data on Amazon.com sales ranks and weeks on sale. Unfortunately, this is a measure of transactions completed prior to the date when we collected data, whereas authors probably care primarily about expected total sales. To address this potential problem, we exclude books that have been on sale for less than one year. We examined the 5229 books published by xLibris that had been

[36] Chevalier and Goolsbee also provide parameter values for the imputation of sales from Amazon rank. However because they do not provide standard errors we would not be able to calculate standard errors for our results that properly account for the imprecision of the parameters used to impute sales.

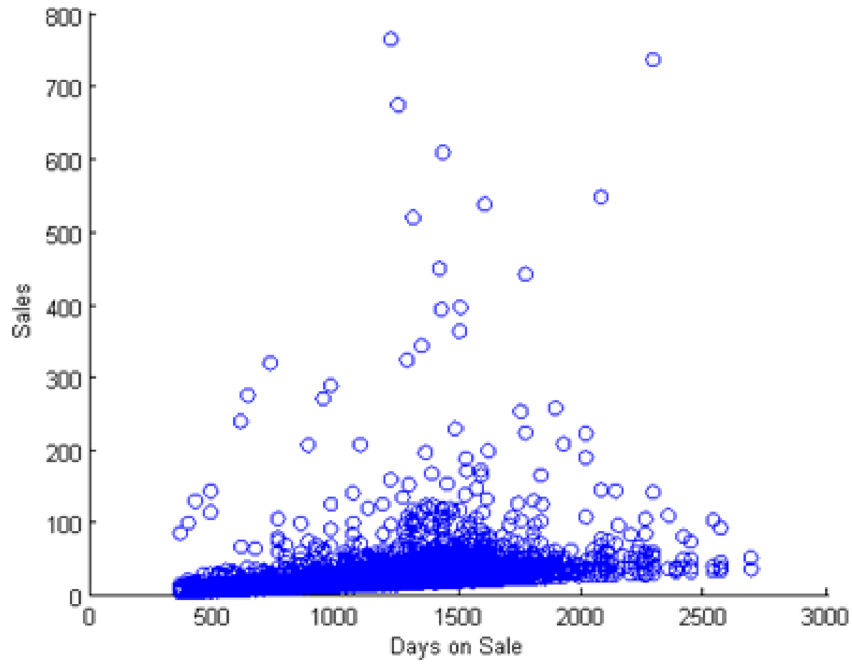

Figure 9: Sales as a function of days on sale

on sale for at least one year (see Figure 3), and found no clear relationship between either the mean or the variance of sales and the period of time over which the book has been available for. We take this as a strong basis under which to proceed under the assumption that all sales take place in the first year (or at least the vast majority). Table 1 presents the summary statistics for our sample along with the summary statistics for BHS. These statistics highlight another potential source of error in our analysis, stemming from the different support of observed sales ranks. Ideally, we'd like our data to cover a similar distribution of observed sales ranks. We rely on the fact that both of these sets of observations capture distributions of obscure titles, with relatively high sales ranks. Some of these differences may reflect the time that elapsed between BHS's findings and when our data was collected (approximately 2 years).

Variable	Obs	Mean	S.D.	Min	Max
BHS(Sales Rank)	861	31,532.85	58,350.92	238	961,367
xLibris(Sales Rank)	5,229	104,210.61	244,587.31	12,968	1,514,226

Table 1. Summary Statistics for BHS and xLibris Data

2.6 Results

We simulate the publishing decision for n = 10000 authors who choose either POD, traditional publishing, or not to publish. xLibris charges $500 to prepare a manuscript for publication and make it available for sale on the internet. We assume the book's price will be $25 .[37] If an author selects traditional publishing, they must also select a quantity cost pair from the set such that

$$(Q, C_{Tra}(Q)) \in \{(200, \$1530), (500, \$3273), (1000, \$3782)\}.^{38}$$

These traditional current prices probably reflect the competition from POD, thus in the counterfactual (the absence of POD) one would imagine that these prices would be higher. For this reason, the welfare improvement from the existence of POD will be underestimated, as even fewer authors would publish in the absence of POD.

We estimated parameter values as discussed above and reported them in the Appendix. We report two sets of results. The results in Table A.1 and A.2 reflect only sampling error (that is, the simulated standard errors do not reflect the possibility that BHS's estimates are not exact), while the results in Table A.3 and A.4 reflect both the simulation error and the standard error in BHS's estimates.

Figure 4 shows the CDF resulting from one of the simulation runs used for Table A. There are more titles with less that 20 copies sold in the simulation than are actually observed, but the fit seems to be good overall. Once we have fitted parameter values,

[37] The approximate average price in our data.

[38] In our simulation we allow an author to print any quantity greater than 1000 for the same average cost (3.782 per copy) to model the notion that an author can always order multiple print runs. We adapt cost data from the print-broker Rjcom.com.

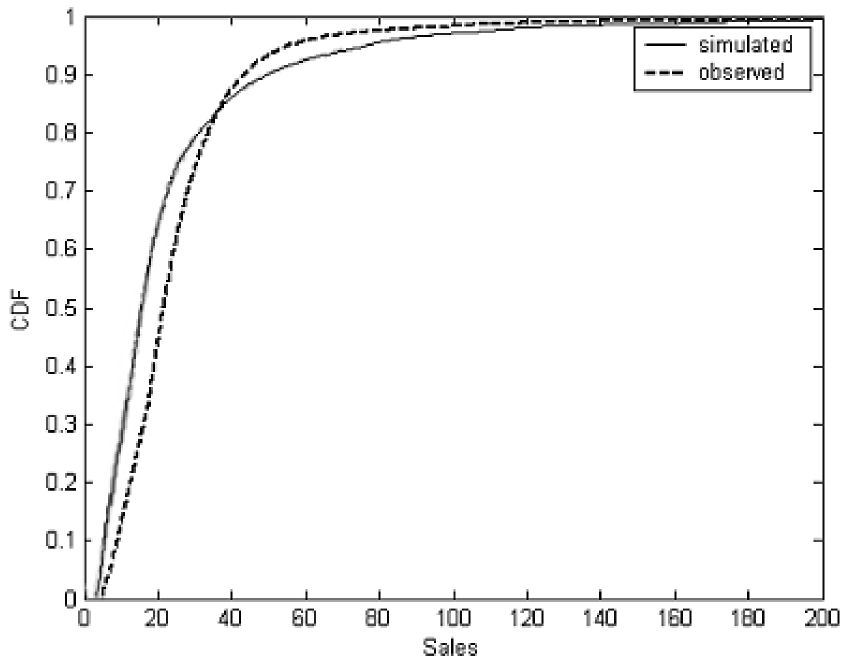

Figure 10: The CDF from a simulation run.

evaluating the consequences of counter-factuals is not difficult. We simply take a number of random draws equal to the estimated number of potential authors and then vary the publishing technologies available to potential authors. In particular, we calculate authors' welfare and the number of titles that would have been published in the absence of POD technology. The results we report in Tables C and D of the Appendix correspond to the parameter values of Tables A and B respectively. Of the 5229 xLibris titles in our sample (all using POD), we estimate only 181 would have been published with traditional technology and authors' welfare would drop from \$1.2 million to \$47,000 (where the outside option, not publishing, is normalized to \$0).

While the estimated welfare effects may not be significant in economic terms, the utility gained by authors in significant: POD technology is responsible for an additional 8300 titles being published and a \$2.2 million increase in author welfare, for an average of approximately \$300 in author surplus per title published over and above the \$500/title publisher's fee. Our results show that while allowing authors access to POD technology does not significantly increase an author's expected profit, however, the estimated utility gained from being published is significant. Therefore, even if POD does not significantly increase a vanity author's expected sales; it does increase welfare by reducing barriers to entry, thereby allowing more authors to indulge in vanity.

POD allows marginal authors to enter the market by reducing the fixed cost of entry and allowing them to capture utility from vanity. Furthermore, the estimated utility gains from vanity for the average author are large relative to their share of expected sales. This seems reasonable, given that print on demand is a relatively new publishing option. Over time, we would expect the fixed costs associated with using POD to fall as the industry is becomes more competitive, limiting a publisher's ability to extract an author's welfare gains from vanity. Further analysis in this regard would require more data on the POD publisher's costs.

2.7 Conclusions

Creativity and art are important in the sense that they have positive externalities. That's why we teach them in our schools and use our taxes to subsidize them. If there are more artists, then the artistic community will be larger, as will the positive externality from the exchange of ideas and techniques. Even for those who do not derive any pecuniary benefit, the production of creative content improves communication skills and enhances mental flexibility. Thus, we are concerned with the effect of technological change not just on the flow of profits, but also on the entry decision.

The manuscripts these authors have written are not viable for trade publication. They are necessarily marginal producers and we should expect their surplus to be small, whatever the available technologies. That POD technology, specifically the reduction in entry costs, should cause a an economically insignificant increase in producer surplus is not surprising because it is only large relative to the small initial producer surplus.

While the POD industry is a small fraction of the entire book industry and thus of limited economic import, the number of titles per year is increasing. Approximately 1% of the books on Amazon are from POD publishers, but all POD titles have been published in the last decade so their share in recently published books is higher and over a long enough time horizon content that has not yet been produced will dominate. Long run consumer welfare from creative-content depends crucially on the continual introduction of new products because the most popular products are typically of recent vintage. Because the pool of potential blockbusters is limited to newly created content and it is difficult to identify which previously obscure artist is likely to create a hit, then if there is more content being produced for the tail, we should expect the quality of future hits to improve.

Appendix 1

Lemma 1. For all $\theta \in R$, $e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L})) > 0$ is continuously differentiable in $\sigma_2^*(L)$, $\sigma_2^*(\not{L})$, and θ.

Proof: Within each history of period 2, each patent holder type's expected profits are fixed as indicated by (1) (3) (5) (7).

$$
\begin{aligned}
\Pi^{PH}(L, Out, \theta) &= \theta + F - e_1 \\
\Pi^{PH}(L, In, \theta) &= \theta + \theta\gamma - e_1 \\
\Pi^{PH}(\not{L}, In, \theta) &= \theta(\gamma + 1)\, p\left(e_2^*(\theta \mid In)\right) - e_2^*(\theta \mid In) - e_1 \\
\Pi^{PH}(\not{L}, Out, \theta) &= F + \theta p\left(e_2^*(\theta \mid Out)\right) - e_2^*(\theta \mid Out) - e_1 \\
\text{where } e_2^*(\theta \mid In) &= \underset{e_2 \in R_+}{\operatorname{argmax}} \left\{\theta(\gamma + 1)\, p(e_2) - e_2\right\} \\
\text{where } e_2^*(\theta \mid Out) &= \underset{e_2 \in R_+}{\operatorname{argmax}} \left\{\theta p(e_2) - e_2\right\}
\end{aligned}
$$

Therefore, the first order condition of a patent holder's profit maximization problem is given in (21) by,

$$
p'(e_1) \begin{bmatrix} \Pi^{PH}(L, Out, \theta) - \Pi^{PH}(\not{L}, Out, \theta) + \\ \sigma_2^*(L)\left(\Pi^{PH}(L, In, \theta) - \Pi^{PH}(L, Out, \theta)\right) + \\ \sigma_2^*(\not{L})\left(\Pi^{PH}(\not{L}, Out, \theta) - \Pi^{PH}(\not{L}, In, \theta)\right) \end{bmatrix} \leq 1 \tag{33}
$$

by (22), this condition holds with equality for any positive amount of optimal effort. Therefore, since $p'(e_1)$ is continuously differentiable and $p''(e_1) < 0$ for all $e_1 \in R_+$, the inverse function theorem proves $p'(e_1)$ has an inverse $g(\cdot) = [p'(e_1)]^{-1}$ which is also continuously differentiable over the interval $p'(R_+)$ such that for all z in the domain of $g(\cdot)$, $g'(z) = \frac{1}{p''(g(z))}$.

Recall,

$$
\psi(\theta) = \begin{bmatrix} E^{PH}\left[\Pi_\theta^{PH} \mid L, Out\right] - E^{PH}\left[\Pi_\theta^{PH} \mid \not{L}, Out\right] + \\ \sigma_2^*(L)\left(E^{PH}\left[\Pi_\theta^{PH} \mid L, In\right] - E^{PH}\left[\Pi_\theta^{PH} \mid L, Out\right]\right) + \\ \sigma_2^*(\not{L})\left(E^{PH}\left[\Pi_\theta^{PH} \mid \not{L}, Out\right] - E^{PH}\left[\Pi_\theta^{PH} \mid \not{L}, In\right]\right) \end{bmatrix} \tag{34}
$$

such that the first order condition is written

$$p'(e_1)\psi(\theta) = 1 \tag{35}$$

Therefore the inverse $g(\cdot)$ has domain $\left(\frac{1}{\psi(\theta)}\right)$

$$g(p'(e_1)) = g\left(\frac{1}{\psi(\theta)}\right) \tag{36}$$

Thus I have

$$e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L})) = g\left(\frac{1}{\psi(\theta)}\right) \tag{37}$$

By the inverse function theorem, and my assumption that $p\prime\prime(e_1) < 0$, I have

$$g'\left(\frac{1}{\psi(\theta)}\right) = \frac{1}{p\prime\prime\left(g\left(\frac{1}{\psi(\theta)}\right)\right)} < 0 \tag{38}$$

$$\frac{\partial g\left(\frac{1}{\psi(\theta)}\right)}{\partial \psi(\theta)} = g'\left(\frac{1}{\psi(\theta)}\right)\frac{\partial\left(\frac{1}{\psi(\theta)}\right)}{\partial \psi(\theta)} \tag{39}$$

$$= g'\left(\frac{1}{\psi(\theta)}\right)\frac{\partial}{\partial \psi(\theta)}[\psi(\theta)]^{-1} \tag{40}$$

Therefore, since

$$g'\left(\frac{1}{\psi(\theta)}\right) = \frac{1}{p\prime\prime\left(g\left(\frac{1}{\psi(\theta)}\right)\right)} < 0 \tag{41}$$

$$\frac{\partial g\left(\frac{1}{\psi(\theta)}\right)}{\partial \psi(\theta)} = -g'\left(\frac{1}{\psi(\theta)}\right)[\psi(\theta)]^{-2} \tag{42}$$

$$> 0 \tag{43}$$

Therefore, since $p(e_1)$ is concave, $g\left(\frac{1}{\psi(\theta)}\right)$ is an increasing function of $\psi(\theta) > 0$. Therefore, the sign of the derivative of optimal effort $g\left(\frac{1}{\psi(\theta)}\right)$ with respect to variable

$x \in \{\theta, \sigma_2^* (L), \sigma_2^* (\not{L})\}$ is such that

$$\frac{\partial g\left(\frac{1}{\psi(\theta)}\right)}{\partial x} > 0 \text{ if } \frac{\partial}{\partial x}[\psi(\theta)] > 0 \tag{44}$$

$$\frac{\partial g\left(\frac{1}{\psi(\theta)}\right)}{\partial x} < 0 \text{ if } \frac{\partial}{\partial x}[\psi(\theta)] < 0 \tag{45}$$

Furthermore I find,

$$\frac{\partial \psi(\theta)}{\partial \sigma_2^*(L)} = [\theta\gamma - F] \tag{46}$$

$$\frac{\partial \psi(\theta)}{\partial \sigma_2^*(\not{L})} = [F + \theta p(e_2^*(\theta \mid Out)) - e_2^*(\theta \mid Out) - \theta(\gamma+1)p(e_2^*(\theta \mid In)) + e_2^*(\theta \mid In)] \tag{47}$$

Using the envelope theorem, because of maximizers $e_2^*(\theta \mid In)$ and $e_2^*(\theta \mid Out)$, I find

$$\frac{\partial \psi(\theta)}{\partial \theta} = [1 + \sigma_2^*(L)\gamma + \sigma_2^*(\not{L})(e_2^*(\theta \mid Out) - (\gamma+1)p(e_2^*(\theta \mid In))) - e_2^*(\theta \mid Out)] \tag{48}$$

Lemma 2: As $\sigma_2^*(\not{L})$ increases, the change in first period patent holder effort is a weakly decreasing function of patent holder type θ, such that

$$\frac{\partial^2 e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L}))}{\partial \sigma_2^*(\not{L})\,\partial\theta} \leq 0, \forall\theta \tag{49}$$

Proof: From Lemma 1, I know that $\frac{\partial e_1^*\left(\theta \mid \sigma_2^*(L), \sigma_2^*\left(\not{L}\right)\right)}{\partial \sigma_2^*\left(\not{L}\right)} < 0$ if and only if $\frac{\partial \psi(\theta)}{\partial \sigma_2^*\left(\not{L}\right)} < 0$. Furthermore, from Lemma 1

$$\frac{\partial \psi(\theta)}{\partial \sigma_2^*(\not{L})} = \tag{50}$$

$$[F + \theta p(e_2^*(\theta \mid Out)) - e_2^*(\theta \mid Out) - \theta(\gamma+1)p(e_2^*(\theta \mid In)) + e_2^*(\theta \mid In)] \tag{51}$$

Therefore, using the envelope theorem because of second period optimal efforts for subgames 3 and 4 ($e_2^* (\theta \mid Out)$ and $e_2^* (\theta \mid In)$) given in (5) and (7), I conclude that

$$\frac{\partial \psi (\theta)}{\partial \sigma_2^* (\not{L}) \partial \theta} = \tag{52}$$

$$\frac{\partial \left[F + \theta p \left(e_2^* (\theta \mid Out)\right) - e_2^* (\theta \mid Out) - \theta (\gamma + 1) p \left(e_2^* (\theta \mid In)\right) + e_2^* (\theta \mid In)\right]}{\partial \theta} \tag{53}$$

$$= p \left(e_2^* (\theta \mid Out)\right) - (\gamma + 1) p \left(e_2^* (\theta \mid In)\right)) \tag{54}$$

$$< 0, \forall \theta > 0 \tag{55}$$

$$= 0, \ \forall \theta \leq 0 \tag{56}$$

Therefore,

$$\frac{\partial^2 \psi (\theta)}{\partial \sigma_2^* (\not{L}) \partial \theta} < 0, \ \forall \theta > 0 \tag{57}$$

$$\frac{\partial^2 \psi (\theta)}{\partial \sigma_2^* (\not{L}) \partial \theta} = 0, \ \forall \theta \leq 0 \tag{58}$$

This proves

$$\frac{\partial^2 e_1^* \left(\theta \mid \sigma_2^* (L), \sigma_2^* (\not{L})\right)}{\partial \sigma_2^* (\not{L}) \partial \theta} \leq 0, \ \forall \theta > 0 \tag{59}$$

Lemma 3. As $\sigma_2^* (L)$ increases, the change in equilibrium first period effort is a weakly increasing function of type θ.

$$\frac{\partial^2 e_1^* \left(\theta \mid \sigma_2^* (L), \sigma_2^* (\not{L})\right)}{\partial \sigma_2^* (L) \partial \theta} \geq 0, \ \forall \theta \tag{60}$$

Proof: From Lemma 1, I know that $\frac{\partial e_1^* \left(\theta \mid \sigma_2^* (L), \sigma_2^* (\not{L})\right)}{\partial \sigma_2^* (L)} > 0$ if and only if $\frac{\partial \psi(\theta)}{\partial \sigma_2^*(L)} > 0$. Furthermore, from Lemma 1,

$$\frac{\partial \psi (\theta)}{\partial \sigma_2^* (L)} = \theta \gamma + F \tag{61}$$

Therefore,

$$\frac{\partial^2 \psi (\theta)}{\partial \sigma_2^* (L) \partial \theta} = \frac{\partial [\theta \gamma + F]}{\partial \theta} \tag{62}$$

$$= \gamma \tag{63}$$

Therefore,

$$\frac{\partial^2 \psi(\theta)}{\partial \sigma_2^*(L)\,\partial\theta} > 0,\ \forall \psi(\theta) > 0 \tag{64}$$

By Lemma 1, this proves

$$\frac{\partial^2 e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L}))}{\partial \sigma_2^*(L)\,\partial\theta} > 0,\ \forall \psi(\theta) > 0 \tag{65}$$

Therefore,

$$\frac{\partial^2 e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L}))}{\partial \sigma_2^*(L)\,\partial\theta} > 0,\ \forall \psi(\theta) > 0 \tag{66}$$

$$\frac{\partial^2 e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L}))}{\partial \sigma_2^*(\not{L})\,\partial\theta} = 0,\ \forall \psi(\theta) < 0 \tag{67}$$

■

Lemma 4: The infringer's expected profits from entry, after observing the existence of litigation in period 1 $(h_1 = L)$, is

(4.1) monotonically increasing in $\sigma_2^*(\not{L})$

(4.2) monotonically decreasing in $\sigma_2^*(L)$

Proof of (4.1): Consider the change in expected profits from entry upon observing the existence of litigation in period 1 $(h_1 = L)$ as I change $\sigma_2^*(\not{L})$.

$$\Pi^{IN}(L, In) \tag{68}$$

Taking the derivative of net expected profit with respect to the infringer's equilibrium entry probability, conditional on observing the absence of litigation I get,

$$\frac{\partial\left[\Pi^{IN}(L, In)\right]}{\partial \sigma_2(\not{L})} \tag{69}$$

$$= \frac{\partial \int [F - \theta\gamma]\,\phi(\theta \mid L)d\theta}{\partial \sigma_2^*(\not{L})} \tag{70}$$

$$= \int [F - \theta\gamma]\,\frac{\partial \phi(\theta \mid L)}{\partial \sigma_2^*(\not{L})} d\theta \tag{71}$$

by Lemma 2 I have

$$\frac{\partial^2 e_1^*\left(\theta \mid \sigma_2^*\left(L\right), \sigma_2^*\left(\not{L}\right)\right)}{\partial \sigma_2^*\left(\not{L}\right) \partial \theta} < 0 \tag{72}$$

In other words, as I increase $\sigma_2^*\left(\not{L}\right)$, any type $\widehat{\theta}$ is reducing its first period effort at a faster rate than any type $\theta < \widehat{\theta}$. Therefore, since

$$\phi\left(\theta \mid L\right) = \frac{p\left(e_1^*\left(\theta\right)\right) \phi_0\left(\theta\right)}{\int\limits_{\theta \in \Theta} p\left(e_1^*\left(\theta\right)\right) \phi_0\left(\theta\right) d\theta} \tag{73}$$

then

$$\frac{\partial \phi(\theta \mid L)}{\partial \sigma_2^*\left(\not{L}\right) \partial \theta} < 0 \tag{74}$$

Therefore,

$$\frac{\partial \Pi^{IN}\left(L, In\right)}{\partial \sigma_2\left(\not{L}\right)} > 0 \tag{75}$$

In short, as $\sigma_2^*\left(\not{L}\right)$ increases, the existence of litigation in period 1 will become less threatening, because in relative terms, lower types will increase the probability that they create litigation faster than higher types as $\sigma_2^*\left(\not{L}\right)$ increases.

Proof of (4.2): Consider the change in expected profits from entry upon observing the existence of litigation in period 1 ($h_1 = L$) as I change $\sigma_2^*\left(L\right)$.

.

Taking the derivative of net expected profit with respect to the infringer's equilibrium entry probability, conditional on observing the existence of litigation I get,

$$\frac{\partial \Pi^{IN}\left(L, In\right)}{\partial \sigma_2\left(L\right)} \tag{76}$$

$$= \frac{\partial \int \left[F - \theta\gamma\right] \phi(\theta \mid L) d\theta}{\partial \sigma_2^*\left(L\right)} \tag{77}$$

$$= \int \left[F - \theta\gamma\right] \frac{\partial \phi(\theta \mid L)}{\partial \sigma_2^*\left(L\right)} d\theta \tag{78}$$

Recall that by Lemma 3 I know

$$\frac{\partial^2 e_1^*\left(\theta \mid \sigma_2^*\left(L\right), \sigma_2^*\left(\not{L}\right)\right)}{\partial \sigma_2^*\left(L\right) \partial \theta} > 0 \tag{79}$$

and therefore

$$\frac{\partial^2 \phi(\theta \mid L)}{\partial \sigma_2^*(L)\,\partial \theta} > 0 \tag{80}$$

$$\frac{\partial \Pi^{IN}(L, In)}{\partial \sigma_2(L)} < 0 \tag{81}$$

■

Lemma 5: The infringer's expected profits from entry after observing the absence of litigation ($h_1 = \not{L}$) in period 1, is

(5.1) monotonically decreasing in $\sigma_2^*(\not{L})$

(5.2) monotonically increasing in $\sigma_2^*(L)$.

Proof of (5.1): Consider the change in the infringer's expected profits from entry upon observing the absence of litigation in period 1 ($h_1 = \not{L}$) as I change $\sigma_2^*(\not{L})$

Taking the derivative of expected profit from entry with respect to the infringer's equilibrium entry probability, conditional on observing the absence of litigation I get,

$$\frac{\partial \Pi^{IN}(\not{L}, In)}{\partial \sigma_2(\not{L})} \tag{82}$$

$$= \frac{\partial \int \left[F + \theta p\left(e_2^*(\theta \mid Out)\right) - \theta \gamma p\left(e_2^*(\theta \mid In)\right)\right] \phi(\theta \mid \not{L}) d\theta}{\partial \sigma_2^*(\not{L})} \tag{83}$$

$$= \int \left[F + \theta p\left(e_2^*(\theta \mid Out)\right) - \theta \gamma p\left(e_2^*(\theta \mid In)\right)\right] \frac{\partial \phi(\theta \mid \not{L})}{\partial \sigma_2^*(\not{L})} d\theta \tag{84}$$

where

$$\begin{aligned} \text{where } e_2^*(\theta \mid In) &= \underset{e_2 \in R_+}{\operatorname{argmax}} \left\{\theta(\gamma+1)\, p(e_2) - e_2\right\} \\ \text{where } e_2^*(\theta \mid Out) &= \underset{e_2 \in R_+}{\operatorname{argmax}} \left\{\theta p(e_2) - e_2\right\} \end{aligned}$$

By Lemma 2 I have

$$\frac{\partial e_1^*\left(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L})\right)}{\partial \sigma_2^*(\not{L})\,\partial \theta} < 0 \tag{85}$$

In other words, as I increase $\sigma_2^*(\not{L})$, any type $\widehat{\theta}$ is reducing its first period effort at a faster rate than any type $\theta < \widehat{\theta}$. Therefore, since

$$\phi\left(\theta \mid \not{L}\right) = \frac{\left(1 - p\left(e_1^*\left(\theta\right)\right)\right)\phi_0\left(\theta\right)}{\int\limits_{\theta\in\Theta}\left(1 - p\left(e_1^*\left(\theta\right)\right)\right)\phi_0\left(\theta\right)d\theta} \tag{86}$$

Lemma 2 proves that

$$\frac{\partial^2\phi(\theta \mid \not{L})}{\partial\sigma_2^*\left(\not{L}\right)\partial\theta} > 0 \tag{87}$$

Therefore,

$$\frac{\partial\Pi^{IN}\left(\not{L}, In\right)}{\partial\sigma_2\left(\not{L}\right)} < 0 \tag{88}$$

In short, as $\sigma_2^*\left(\not{L}\right)$ increases, the absence of litigation in period 1 will become more threatening, because in relative terms, lower types will increase the probability that they create litigation faster than higher types as $\sigma_2^*\left(\not{L}\right)$ increases

Proof of (5.2): Consider the change in the infringer's expected profits from entry upon observing the existence of litigation in period 1 $(h_1 = L)$ as I change $\sigma_2^*\left(L\right)$.

$$\frac{\partial\Pi^{IN}\left(\not{L}, In\right)}{\partial\sigma_2\left(L\right)} \tag{89}$$

$$= \frac{\partial\int\left[F + \theta p\left(e_2^*\left(\theta \mid Out\right)\right) - \theta\gamma p\left(e_2^*\left(\theta \mid In\right)\right)\right]\phi(\theta \mid \not{L})d\theta}{\partial\sigma_2^*\left(L\right)} \tag{90}$$

$$= \int\left[F + \theta p\left(e_2^*\left(\theta \mid Out\right)\right) - \theta\gamma p\left(e_2^*\left(\theta \mid In\right)\right)\right]\frac{\partial\phi(\theta \mid \not{L})}{\partial\sigma_2^*\left(L\right)}d\theta \tag{91}$$

where

$$\begin{aligned}\text{where } e_2^*\left(\theta \mid In\right) &= \underset{e_2\in R_+}{\operatorname{argmax}}\left\{\theta\left(\gamma+1\right)p\left(e_2\right) - e_2\right\}\\ \text{where } e_2^*\left(\theta \mid Out\right) &= \underset{e_2\in R_+}{\operatorname{argmax}}\left\{\theta p\left(e_2\right) - e_2\right\}\end{aligned}$$

Recall that by Lemma 3 I know

$$\frac{\partial^2 e_1^*\left(\theta \mid \sigma_2^*\left(L\right), \sigma_2^*\left(\not{L}\right)\right)}{\partial\sigma_2^*\left(L\right)\partial\theta} > 0 \tag{92}$$

and therefore

$$\frac{\partial^2 \phi(\theta \mid \not L)}{\partial \sigma_2^* (L) \, \partial \theta} < 0 \tag{93}$$

$$\frac{\partial \Pi^{IN} (\not L, In)}{\partial \sigma_2 (L)} > 0 \tag{94}$$

■

Lemma 6: In any Bayesian equilibrium $\sigma_2^* (L) > 0$ only if $\sigma_2^* (\not L) = 1$.

Proof: By assumption,

$$\left[\Pi^{IN} (\not L, In \mid \theta) - \Pi^{IN} (\not L, Out \mid \theta)\right] > \Pi^{IN} (L, In \mid \theta) - \Pi^{IN} (L, Out \mid \theta) \tag{95}$$

Therefore using (23) and (24) I see that if $\phi(\theta \mid \not L) \geq \phi(\theta \mid L)$ then

$$\Pi^{IN} (\not L, In) > \Pi^{IN} (L, In) . \tag{96}$$

Furthermore, if $\sigma_2^* (\not L) < \sigma_2^* (L)$ then from (28) and (29) I know $\frac{\partial e_1^* \left(\theta \mid \sigma_2^*(L), \sigma_2^* \left(\not L\right)\right)}{\partial \theta} \geq 0 \ \forall \theta$, which implies $\phi(\theta \mid \not L) \geq \phi(\theta \mid L)$ and thus $\Pi^{IN} (\not L, In) > \Pi^{IN} (L, In)$ which contradicts $\sigma_2^* (\not L) < \sigma_2^* (L)$. Therefore I conclude that in any equilibrium $\sigma_2^* (\not L) \geq \sigma_2^* (L)$.

Furthermore, since the inequality in (98) is strict, it cannot be true that

$$\Pi^{IN} (\not L, In) = \Pi^{IN} (L, In) = 0. \tag{97}$$

Therefore, the infringer cannot be indifferent to entry under both histories $\{\not L\}$ and $\{L\}$. Thus $\sigma_2^* (L) > 0$ only if $\sigma_2^* (\not L) = 1$.■

Proposition1: For any setting of the parameters and any continuous function $p(\cdot)$ under my assumptions, and for any possible prior beliefs $\phi_0(\cdot)$ and any type space $\Theta \subset R$, there exists a unique Bayesian equilibrium.

Proof of Existence: I can check to see if there is an equilibrium when the infringer best responses are in pure strategies using the following procedure:

1. Consider a hypothetical best response $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\}$ from the set

$$\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} \in \{\{Out, Out\}, \{Out, In\}, \{In, In\}\},$$

excluding $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{In, Out\}$ because of Lemma 6.

2. Take $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\}$ and plug it into each patent holder's profit maximization problem to find the optimal choice of effort for each patent holder type:

$$\{e_1^*(\theta \mid \hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})) \ \forall\theta\}$$

3. Using $\{e_1^*(\theta \mid \hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})) \ \forall\theta\}$, calculate $\phi(\theta \mid L)$ and $\phi(\theta \mid \not{L})$.

4. Using $\phi(\theta \mid L)$ and $\phi(\theta \mid \not{L})$, calculate $\Pi^{IN}(L, In)$ and $\Pi^{IN}(\not{L}, In)$.

5. Using $\Pi^{IN}(L, In)$ and $\Pi^{IN}(\not{L}, In)$ I know

$$\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L}), \{e_1^*(\theta \mid \hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})) \ \forall\theta\}\}$$

is a Bayesian equilibrium only if, for all $h_1 \in \{L, \not{L}\}$

$$\begin{aligned} \hat{\sigma}_2^*(h_1) &= \{In\}, \quad \text{only if } \Pi^{IN}(h_1, In) > 0 \\ &\quad \text{and} \\ \hat{\sigma}_2^*(h_1) &= \{Out\}, \ \text{only if } \Pi^{IN}(h_1, In) < 0 \end{aligned}$$

because by definition

$$\begin{aligned} &\sigma_2^*(h_1) = 1, && \text{if } \Pi^{IN}(h_1, In) > 0 \\ &\sigma_2^*(h_1) \in (0,1), && \text{if } \Pi^{IN}(h_1, In) = 0 \\ &\sigma_2^*(h_1) = 0, && \text{if } \Pi^{IN}(h_1, In) < 0 \end{aligned}$$

When the potential infringer's best responses are not in pure strategies, I can resume the search for equilibria by asking the following question. Given the optimal first period effort for each type of patent holder when infringer best responses are $\{\sigma_2^*(L), \sigma_2^*(\not{L})\} = \{Out, In\}$, which of the best responses: $\{Out, Out\}$ or $\{In, In\}$, would the infringer wish to deviate too?

Consider the following cases of the deviations to $\{Out, Out\}$ or $\{In, In\}$ given the set of patent holder effort $(e_1^*(\theta \mid Out, In))\,\forall\theta \in \Theta$.

Case 1: Assume that given the set of patent holder effort $\{e_1^*(\theta \mid Out, In)\;\forall\theta \in \Theta\}$, the infringer wishes to play $\{\sigma_2^*(L), \sigma_2^*(\not{L})\} = \{Out, Out\}$. Then by

$$\{\sigma_2^*(L), \sigma_2^*(\not{L})\} \notin \{\{Out, Out\}, \{Out, In\}, \{In, In\}\},$$

I know its expected profits from entry conditional on observing $\{\not{L}\}$ are negative given $\{e_1^*(\theta \mid Out, In)\;\forall\theta \in \Theta\}$, but positive given $\{e_1^*(\theta \mid Out, Out)\;\forall\theta \in \Theta\}$. Thus by Lemma 5, the infringer's second period expected profits from entry, after observing the absence of litigation, are continuous and decreasing in $\sigma_2^*(\not{L})$. Therefore, there exists a set of efforts $\{e_1^*(\theta \mid Out, \sigma_2)\;\forall\theta \in \Theta\}$ for some $\sigma_2 \in (0,1)$, such that the potential infringer expects zero profit from entering upon observing $\{\not{L}\}$. Thus there exists an equilibrium

$$\{e_1^*(\theta \mid Out, \sigma_2)\;\forall\theta \in \Theta\} \text{ and } \{\sigma_2^*(L), \sigma_2^*(\not{L})\} = \{Out, \sigma_2\}$$

for some $\sigma_2 \in (0,1)$.

Case 2: Assume that given the set of patent holder efforts $\{e_1^*(\theta \mid Out, In)\;\forall\theta \in \Theta\}$, the infringer wishes to play $\{\sigma_2^*(L), \sigma_2^*(\not{L})\} = \{In, In\}$. Then by

$$\{\sigma_2^*(L), \sigma_2^*(\not{L})\} \notin \{\{Out, Out\}, \{Out, In\}, \{In, In\}\},$$

I know its expected profits from entry conditional on observing $\{L\}$ are positive given $\{e_1^*(\theta \mid Out, In) \;\; \forall \theta \in \Theta\}$, but negative given $\{e_1^*(\theta \mid In, In) \;\; \forall \theta \in \Theta\}$. Thus by Lemma 4, the infringer's second period expected profits from entry, after observing the absence of litigation, are continuous and decreasing in $\sigma_2^*(L)$. Therefore, there exists a set of efforts $\{e_1^*(\theta \mid \sigma_2, In) \;\; \forall \theta \in \Theta\}$ for some $\sigma_2 \in (0,1)$, such that the potential infringer expects zero profit from entering upon observing $\{L\}$. Thus there exists an equilibrium

$$\{e_1^*(\theta \mid \sigma_2, In) \;\; \forall \theta \in \Theta\} \text{ and } \{\sigma_2^*(L), \sigma_2^*(\not{L})\} = \{\sigma_2, In\}$$

for some $\sigma_2 \in (0,1)$.

Proof of Uniqueness:

Part 1: proving unique equilibrium best response for the infringer

Case PS: (Pure strategy equilibrium): Assume there exists an equilibrium with infringer best responses:

$$\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\}$$

such that

$$\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} \in \{\{Out, Out\}, \{Out, In\}, \{In, In\}\} \tag{98}$$

Case PS.1: It cannot be the case that both $\{Out, Out\}$ $\{Out, In\}$ are in equilibrium because if $\{Out, In\}$ is an equilibrium, by Lemma 5 I know that the infringer's expected profits upon observing the $\{\not{L}\}$, are monotonically increasing as $\sigma_2^*(\not{L}) \to 0$. Therefore, it cannot be the case that case that the infringer expects positive profits from entry upon observing $\{\not{L}\}$ when playing $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{Out, In\}$, but also expect negative profits from entry upon observing $\{\not{L}\}$ when playing $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{Out, Out\}$.

Case PS.2: It cannot be the case that both $\{In, In\}$ $\{Out, In\}$ are in equilibrium because if $\{Out, In\}$ is an equilibrium, by Lemma 4 I know that the infringer's expected profits upon observing $\{L\}$, are monotonically increasing as $\sigma_2^*(L) \to 0$. Therefore, it cannot be the case that case that the infringer expects negative profits from entry upon observing $\{L\}$ when playing $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{Out, In\}$, but also expect positive profits from entry upon observing $\{L\}$ when playing $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{In, In\}$.

Case PS.3: It cannot be the case that both $\{In, In\}$ $\{Out, Out\}$ are in equilibrium because if $\{In, In\}$ is an equilibrium, by Lemma 4 I know that the infringer's expected profits upon observing $\{\not{L}\}$, are monotonically increasing as $\sigma_2^*(L) \to 0$. Therefore, it cannot be the case that case that the infringer expects positive profits from entry upon observing $\{\not{L}\}$ when playing $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{In, In\}$, but also expect positive profits from entry upon observing $\{\not{L}\}$ when playing $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{Out, Out\}$.

Case MS: Mixed Strategy Best responses:

$$\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{\{\hat{\sigma}_2^*, In\}, \{Out, \check{\sigma}_2^*\}\} \tag{99}$$

for some $\hat{\sigma}_2^*$ and $\check{\sigma}_2^* \in (0,1)$

It cannot be the case that $\{\{\hat{\sigma}_2^*, In\}, \{Out, \check{\sigma}_2^*\}\}$ are both in equilibrium for any $\hat{\sigma}_2^*$ and $\check{\sigma}_2^* \in (0,1)$, because if $\{\hat{\sigma}_2^*, In\}$ is in equilibrium, it must be the case that the infringer expects positive entry profits upon observing $\{L\}$ when playing $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{Out, In\}$ and expects negative profits from entry upon observing $\{L\}$ when playing $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{In, In\}$. Therefore since by Lemma 4 the infringer's expected profits from entry, upon observing $\{\not{L}\}$, are increasing as $\sigma_2^*(\not{L}) \to 0$, it cannot be the case that the infringer is indifferent to entry, upon observing $\{\not{L}\}$, for any $\sigma_2^*(\not{L}) < 1$, when there exists another equilibrium $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\} = \{\hat{\sigma}_2^*, In\}$.

Furthermore, I know that when there is a equilibrium in mixed strategies, there exists only one entry probability $\sigma_2^*(h_1)$ that makes the infringer indifferent upon observing h_1, because of the monotonicity of expected profits shown in Lemmas 4 and 5 (See Case 1 and 2 in the proof of existence above).

Part 2: Uniqueness of Patent Holder Effort:

Each patent holder's choice of effort is uniquely determined for any given best response $\sigma_2^*(L), \sigma_2^*(\not{L})$ (from (21)-(30)). Therefore, since there is only one equilibrium set of best responses $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\}$ (Shown in Part 1 of this proof), there exists only one equilibrium $\{\hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})\}$, $\{e_1^*(\theta \mid \hat{\sigma}_2^*(L), \hat{\sigma}_2^*(\not{L})) \ \forall \theta \in \Theta\}$ such that the infringer is profit maximizing according to their beliefs, and beliefs are consistent with the patent holder's equilibrium actions.■

Proposition 2: There exists a non-empty region of the parameter space, prior beliefs $\phi_0(\theta)$, and type spaces $\Theta \subset R$ such that the equilibrium is reputational.

Proof: The equilibrium is reputational as long as the infringer's action in period two is conditional on the first period outcome with positive probability. Therefore I must find the parameter space that excludes the two equilibrium two trivial equilibria. Recall the infringer's expected profits from (18):

$$\begin{aligned}\Pi^{IN}(\not{L}, In) &= \\ &F - \int \theta(\gamma+1)\, p(e_2^*(\theta \mid In))\, \phi(\theta \mid \not{L})\, d\theta + \int \theta p(e_2^*(\theta \mid Out))\, \phi(\theta \mid \not{L})\, d\theta \\ \Pi^{IN}(L, In) &= F - \int \theta\gamma\phi(\theta \mid L)\, d\theta\end{aligned}$$

Therefore, for any $\Pi^{IN}(\not{L}, In)$, there exists a licensing fee F, such that

$$F < \int \theta\gamma\overline{\phi}(\theta \mid L)\, d\theta \tag{100}$$

$$F > \begin{bmatrix} \int \theta(\gamma+1)\, p(e_2^*(\theta \mid In))\, \underline{\phi}(\theta \mid \not{L})\, d\theta ... \\ -\int \theta p(e_2^*(\theta \mid Out))\, \underline{\phi}(\theta \mid \not{L})\, d\theta \end{bmatrix} \tag{101}$$

$$\text{Where } \overline{\phi}(\theta \mid L) \equiv \frac{p(e_1^*(\theta \mid 1,1))\, \phi_0(\theta)}{\int_{\theta\in\Theta} p(e_1^*(\theta \mid 1,1))\, \phi_0(\theta)\, d\theta} d\theta \tag{102}$$

$$\text{Where } \underline{\phi}(\theta \mid \not{L}) \equiv \frac{(1 - p(e_1^*(\theta \mid 0,0)))\, \phi_0(\theta)}{\int_{\theta\in\Theta} (1 - p(e_1^*(\theta \mid 0,0)))\, \phi_0(\theta)\, d\theta} \tag{103}$$

and that will result in reputational equilibria. From Lemma 6, I know $\Pi^{IN}(\not{L}, In) > \Pi^{IN}(L, In)$, therefore, there always exists an F such that

$$F - \begin{bmatrix} \int \theta(\gamma+1)\, p(e_2^*(\theta \mid In))\, \phi(\theta \mid \not{L})\, d\theta + \\ \int \theta p(e_2^*(\theta \mid Out))\, \phi(\theta \mid \not{L})\, d\theta \end{bmatrix} > 0 \tag{104}$$

$$F - \int \theta\gamma\phi(\theta \mid L)\, d\theta < 0 \tag{105}$$

whenever $\Pi^{IN}(\not L, In) > 0$.■

Proposition 3: In any reputational equilibrium, there exists $\overline{\theta} \in R$ such that first period effort is decreasing in θ if and only if $\theta > \overline{\theta}$.

Proof: From Lemma 1 I know $\frac{\partial e_1^*\left(\theta | \sigma_2^*(L), \sigma_2^*\left(\not L\right)\right)}{\partial \theta} < 0$ if and only if $\frac{\partial \psi(\theta)}{\partial \theta} < 0$. Furthermore, from (50) I know

$$\frac{\partial \psi(\theta)}{\partial \theta} = \tag{106}$$

$$= \left[1 + \sigma_2^*(L)\gamma + \sigma_2^*(\not L)\begin{pmatrix} p\left(e_2^*(\theta \mid Out)\right) + ... \\ -(\gamma+1)\,p\left(e_2^*(\theta \mid In)\right)\end{pmatrix} - p\left(e_2^*(\theta \mid Out)\right)\right] \tag{107}$$

$$= 1 + \sigma_2^*(L)\gamma - \sigma_2^*(\not L)(\gamma+1)\,p\left(e_2^*(\theta \mid In)\right) - \left(1 - \sigma_2^*(\not L)\right)p\left(e_2^*(\theta \mid Out)\right) \tag{108}$$

furthermore, as $\theta \to \infty$, $p\left(e_2^*(\theta \mid In)\right), p\left(e_2^*(\theta \mid Out)\right) \to 1$. Therefore,

$$\lim_{\theta \to \infty} \frac{\partial \psi(\theta)}{\partial \theta} = 1 + \sigma_2^*(L)\gamma - \sigma_2^*(\not L)(\gamma+1) - \left(1 - \sigma_2^*(\not L)\right) \tag{109}$$

$$= \gamma\left(\sigma_2^*(L) - \sigma_2^*(\not L)\right) \tag{110}$$

From Lemma 6, I know $\sigma_2^*(\not L) > \sigma_2^*(L)$ for all reputational equilibria. Therefore $\sigma_2^*(L) - \sigma_2^*(\not L) < 0$, and γ strictly positive implies.

$$\lim_{\theta \to \infty} \frac{\partial \psi(\theta)}{\partial \theta} = \gamma\left(\sigma_2^*(L) - \sigma_2^*(\not L)\right) < 0 \tag{111}$$

therefore there exists $\overline{\theta} \in R$ such that for all $\theta > \overline{\theta}$, $\frac{\partial e_1^*\left(\theta | \sigma_2^*(L), \sigma_2^*\left(\not L\right)\right)}{\partial \theta} < 0$.■

Proposition 4: In any reputational equilibrium, there exists $\overline{\overline{\theta}} \in R$ and $\underline{\theta} \in R$ such that

$e_1^*(\theta) = 0, \forall \left\{\theta \mid \theta > \overline{\overline{\theta}} \text{ or } \theta < \underline{\theta}\right\}$.

Proof: From Proposition 3, I know that as $\theta \to \infty$

$$\frac{\partial \psi(\theta)}{\partial \theta} \to \gamma\left(\sigma_2^*(L) - \sigma_2^*(\not L)\right) < 0 \tag{112}$$

Therefore, there exists a $\overline{\overline{\theta}} \in R$ such that for all $\theta > \overline{\overline{\theta}}$, $\psi(\theta) = 0$, and therefore for all types $\theta > \overline{\overline{\theta}}$, $e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L})) = 0$.

Furthermore, as $\theta \to -\infty$, then $p(e_2^*(\theta \mid In)), p(e_2^*(\theta \mid Out)) \to 0$. Therefore,

$$\lim_{\theta \to -\infty} \psi(\theta) = -\infty \tag{113}$$

so there exists $\underline{\theta} \in R$ such that for all types $\theta < \underline{\theta}$, $\psi(\theta) = 0$ and therefore for all types $\theta < \underline{\theta}$, $e_1^*(\theta \mid \sigma_2^*(L), \sigma_2^*(\not{L})) = 0$.■

Appendix 2

Appendix 2.A: Tables

	Parameters					
	α	β	μ_v	σ_v	σ_ε	Potential Authors
XLibris	-0.71	-0.065	6.01	0.52	1.09	12,400
$\overline{s}^2_{\hat{\theta},\mathrm{Sim}}$	(0.02)	(0.004)	(0.49)	(0.29)	(0.03)	(1,100)
$s^2_{\hat{\theta},Boot}$	(0.13)	(0.019)	(1.18)	(0.42)	(0.17)	(5,300)

Table 2 (w/o stage 1 BHS variation), T=100 simulation runs

	Parameters					
	α	β	μ_v	σ_v	σ_ε	Potential Authors
XLibris	-0.64	-0.101	5.98	0.48	1.04	22,000
$\overline{s}^2_{\hat{\theta},\mathrm{Sim}}$	(0.25)	(0.065)	(0.42)	(0.49)	(0.11)	(7,300)
$s^2_{\hat{\theta},Boot}$	(0.15)	(0.230)	(0.98)	(0.31)	(0.16)	(13,000)

Table 3 (w/ stage 1 BHS variation), T = 100 simulation runs

	Authors Welfare	Titles Published
xLibris		
POD	1,240,000 (540,000)	5229 (-)
Traditional	47,000 (84,000)	181 (144)
xLibris Change	-1,193,000	5048
Market Change	-2,220,000	8390

Table 4 (w/o stage 1 BHS variation), T = 100 simulation runs

	Authors Welfare	Titles Published
xLibris		
POD	2,100,000 (1,100,000)	5229 (-)
Traditional	1,700,000 (1,100,000)	514 (1200)
xLibris Change	-1,400,000	4715
Market Change	-2,600,000	8770

Table 5 (w/ stage 1 BHS variation), T = 100 simulation runs

Sales Stats	w/o stage 1 BHS variation	w stage 1 BHS variation
Mean	14.7	15.1
S.D(of Mean)	1.36	2.43

Table 6 (Summary Statistics for the $B = 1000$ bootstrap samples of size $W = 1250$)

Appendix 2.B: Identification

In this section we demonstrate how the CDF of simulated POD sales is affected by changes in the estimated parameter values. In other words working through the intuition of the claim that for some $\epsilon_j > 0$ and all $j \in \{1, 2, 3, 4, 5\}$ we have

$$\left\{Z(Sales_i \mid \widehat{\theta}_j + \epsilon_j, \widehat{\theta}_{-j})\right\}_{i=1}^{5229} \neq \left\{Z(Sales_i \mid \widehat{\theta}_j, \widehat{\theta}_{-j} + \epsilon_{-j})\right\}_{i=1}^{5229}, \forall\ \epsilon_{-j} \in R^4.$$

We will begin by showing how all variables have independent effects when assuming $\sigma_\varepsilon = 0$. Once this is done we can demonstrate how our choice of σ_ε has a unique impact on the CDF of simulated POD sales.

Throughout we make the distinction between the CDF of sales for the entire population of authors, and the CDF of POD sales. While the distribution of all authors sales will have an impact on the resulting CDF of POD sales, we are exclusively concerned with matching the CDF of POD sales to the observed CDF of POD sales in our data.

The Distribution of Expected Sales.

The value of α determines the minimum value of CDF of sales in the population(all authors before publishing decisions). This will then govern the x intercept of the CDF of POD author sales, since α determines the minimum value of expected sales for all authors. Figure 5 shows the impact of a change in α.

The value of β determines the acceleration of the CDF of expected sales towards 1. If $\beta = 0$, the CDF converges immediately to 1. As we increase β, the CDF converges more slowly along the x axis to 1. Figure 6 shows the impact of a change in β on the CDF of POD sales.

However, we need to fit the distribution of POD sales, not the CDF of all authors sales. Therefore, we must consider the other parameters impact. The distribution of expected sales of all authors is simply determined by the parameters governing the Pareto distribution. In contrast, distribution of POD sales is determined by the conditional distribution of expected sales, given that the author chooses POD. In a sense, we only observe a portion of authors contained within a subspace of author types. The following thresholds delineate the subspace of authors types that will choose POD.

$$\underline{s}(v) = \frac{C_{POD} - v}{R_{POD}} \tag{B.1}$$

$$\overline{s}(\theta) = \inf\left\{s \mid \mu^E(POD \mid t, \theta) - \mu^E(TRD_{Q_i} \mid t, \theta) > 0,\ \forall Q_i \in Q\right\}^{39} \tag{B.2}$$

Figure 3 shows these boundaries in the type space $S \times V$.

The Distribution of Vanity

The values of μ_v and σ_v, determine the intercept and the shape of the CDF of POD sales at the lower end of the distribution. Consider $\mu_v = k$, and $\sigma_v = 0$: only authors with expected sales expectation greater than some threshold

$$\underline{s}(k) = \frac{C_{POD} - k}{R_{POD}} \tag{B.3}$$

will publish using POD. Our depiction of type choices would then become a single horizontal slice of Figure 3 at $v = k$. Therefore, in the absence of any noise, and with $\sigma_v = 0$, this will define the x intercept of the CDF of observed POD sales such that:

$$\inf\{s \mid CDF(s) > 0\} = \min\{\underline{s}(k), e^{\alpha}\} \tag{B.4}$$

Therefore, as $\mu_v \rightarrow \infty$, $\underline{s}(k) \rightarrow e^{\alpha}$, thus the lower end of the CDF of POD sales will have the same intercept as the population CDF (because with all potential authors having the same vanity, and that vanity being sufficiently large, no potential author will choose not to publish). As we decrease μ_v we will eventually increase the x intercept, and flatten the CDF, because removing the lower end of sales expectations will make the densities of observed POD sales more uniform (as we only observe an upper potion of the Pareto distribution of expected sales, as low sales expectation authors choose not to publish). Figure 7 illustrates this effect below.

[39] Note that this threshold is independent of an authors vanity. The intution being that for resonable approxiation of the CDF of POD sales, an author deciding between POD and traditional publishing already has sufficient vanity to publish. In other words, the upper threshold must be above the lower threshold, or no POD authors will be observed.

Now consider an increase in σ_v. Now some authors with lower sales expectations, but higher vanity will be drawn into publishing with POD. This will create a inward bend in the lower portion of the CDF and push a small portion of the CDF below $\underline{s}(k)$ defined by equation B.3. Figure 8 illustrates this effect below.

Notice that without noise, nothing will change the point where the CDF of observed sales hits 100%, because without noise, the upper limit of observed sales is simply determined by the unique point in expected sales $\overline{s}$ such that

$$p\overline{s}R_{Tra} - C_{200} = p\overline{s}R_{POD} - C_{POD}$$

$$\overline{s} = \frac{C_{200} - C_{POD}}{p\left(R_{Tra} - R_{POD}\right)}$$

The Variance of Noise: σ_ε

The value of σ_ε, determines the cutoff of the upper end of simulated sales: This is true because as we increase σ_ε, the level of expected sales above which an author will choose traditional printing ($\overline{s}(\sigma_\varepsilon)$) also increases, because an increase in σ_ε increases the expected profit of POD sales relative to traditional publishing at all sales levels. This will move the upper threshold $\overline{s}$ to the right. Consequently this change will shift the point where the CDF of simulated POD sales gets close to 1 to the right. No other change in the other parameters can have any impact on the upper bound of the CDF of POD sales. Figure 3 illustrates the impact of a change in σ_ε on the boundaries in the type space delineating optimal author choice. Figure 9 illustrates the impact on the CDF of POD sales.

Potential Author Population Estimate $\Gamma(\widehat{\theta})$

Fitting the CDF's of simulated to observed sales does not require the number of potential authors (it does not appear in our objective function). For each simulation run t, the estimate of the parameters $\widehat{\theta}_t$ results in a certain number of the n =10,000 potential author simulants choosing to publish with POD. Let $\eta(\widehat{\theta}_t)$ be the number of simulants that choose to publish with POD, and N be the number of POD authors in the data.

The estimated probability that a random potential author i drawn from the distributions parameterized by $\widehat{\theta}_t$ then becomes

$$\text{Prob}\left(a^*(s_i, v_i \mid \widehat{\theta}_t) = POD\right) = \frac{\eta(\widehat{\theta}_t)}{n}.$$

Therefore, our estimate of the potential author population, denoted $\Gamma(\widehat{\theta}_t)$, is

$$\Gamma(\widehat{\theta}_t) = \frac{n}{\eta(\widehat{\theta}_t)} N.$$

Consider an example using the data from xLibris with a sample of $N = 5229$ POD authors. If we find that the number of simulants that choose POD, once we have fitted parameter values $\widehat{\theta}$, is $\eta(\widehat{\theta}_t) = 3243$, then our potential author population estimate is

$$\Gamma(\widehat{\theta}_t) = \frac{10,000}{3243} 5229 = 16,123.$$

The Distribution of Vanity

In this section we discuss the distributional assumptions of vanity, and the implications of other potential assumptions. Consider table C, which shows the results from using a normal distribution of vanity.

	Parameters					
	α	β	μ_v	σ_v	σ_ε	Potential Authors
xLibris	-0.58 (0.83)	-0.074 (0.066)	471(10.7)	155(76)	.98 (0.29)	31,300 (24,900)

Table 7 (Normally distributed Vanity w/o stage 1 BHS variation), T=10 simulation runs

	Parameters					
Mass	α	β	μ_v	σ_v	σ_ε	Potential Authors
$\gamma = .1$	-0.61(0.05)	-0.088(0.006)	5.91(0.33)	0.57(0.05)	1.00(0.02)	13,000(9,100)
$\gamma = .5$	-0.81(0.12)	-0.070(0.023)	6.32 (0.34)	0.63(0.09)	1.15(0.11)	29,000(14,000)
$\gamma = .9$	-0.72(0.39)	-0.091(0.007)	6.37(0.51)	0.63(0.23)	1.02(0.15)	130,000(108,000)

Table 8 (Log Normal Vanity with Mass Point γ w/o stage 1 BHS variation), T=100

	$\alpha = 0$	β	μ_v	σ_v	σ_ε	Potential Authors
XLibris	0.0(-)	-0.080(0.009)	5.38(0.41)	0.41(0.20)	1.29(0.13)	39,400(21,000)

Table 9 (Log Normal Vanity and $\alpha = 0$ w/o stage 1 BHS variation), T=100

One consequence of using the normal distribution is that the estimate of potential authors increases. Because the normal distribution is symmetric around the mean, a larger fraction of simulated potential authors choose not to publish, because they had insufficient vanities to justify the fixed costs of POD.

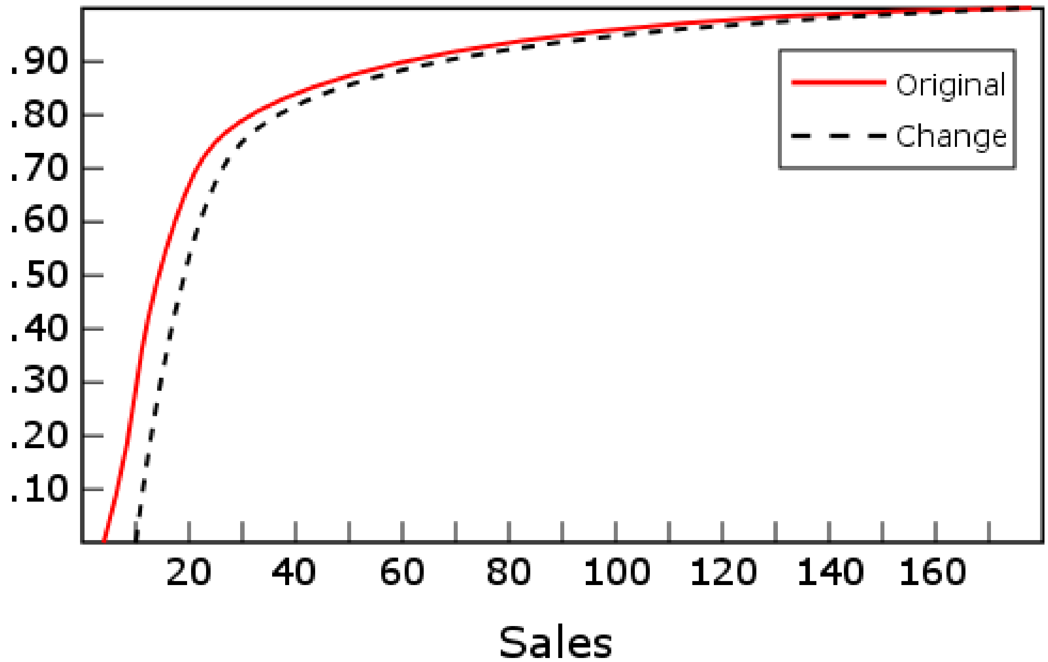

Figure 11: Change in the CDF of simulated POD sales for an increase in α.

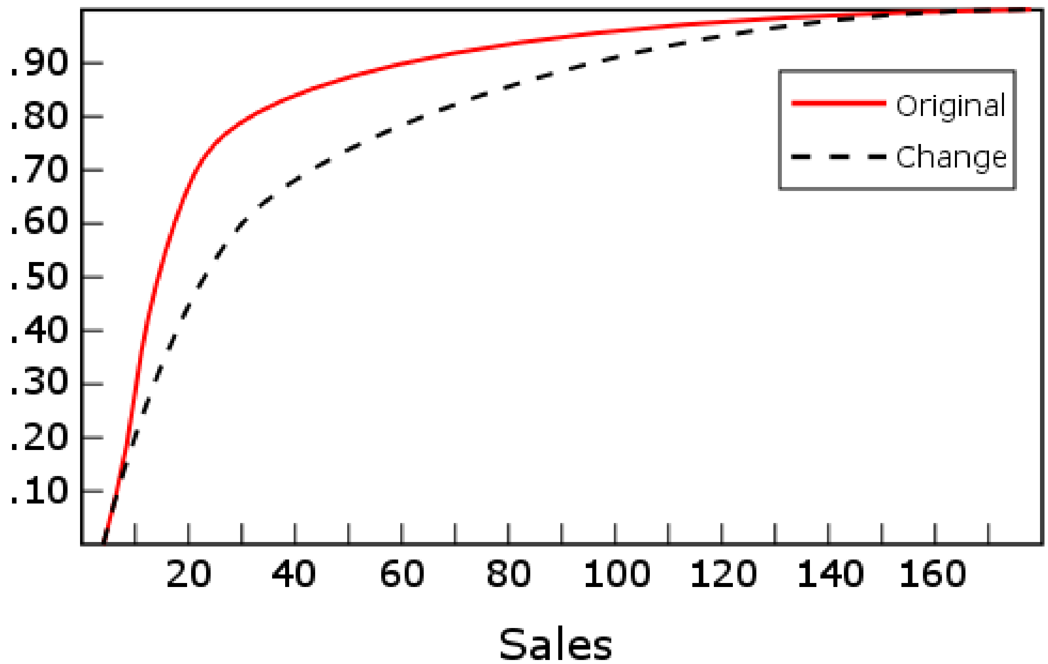

Figure 12: Change in the CDF of simulated POD sales for an increase in β.

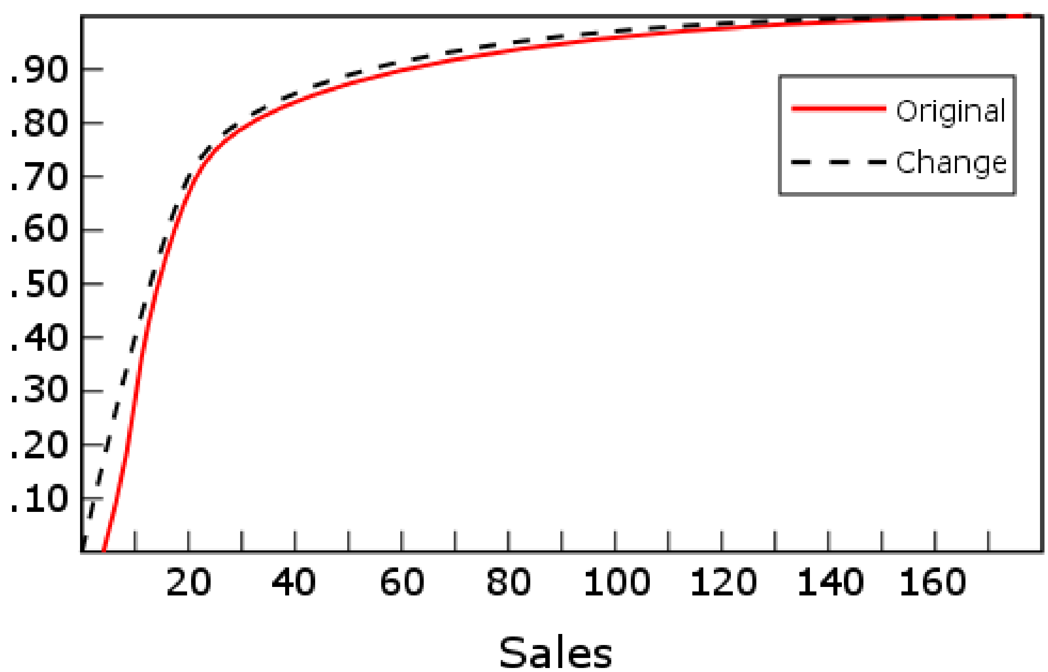

Figure 13: Change in the CDF of simulated POD sales for an increase in μ_v.

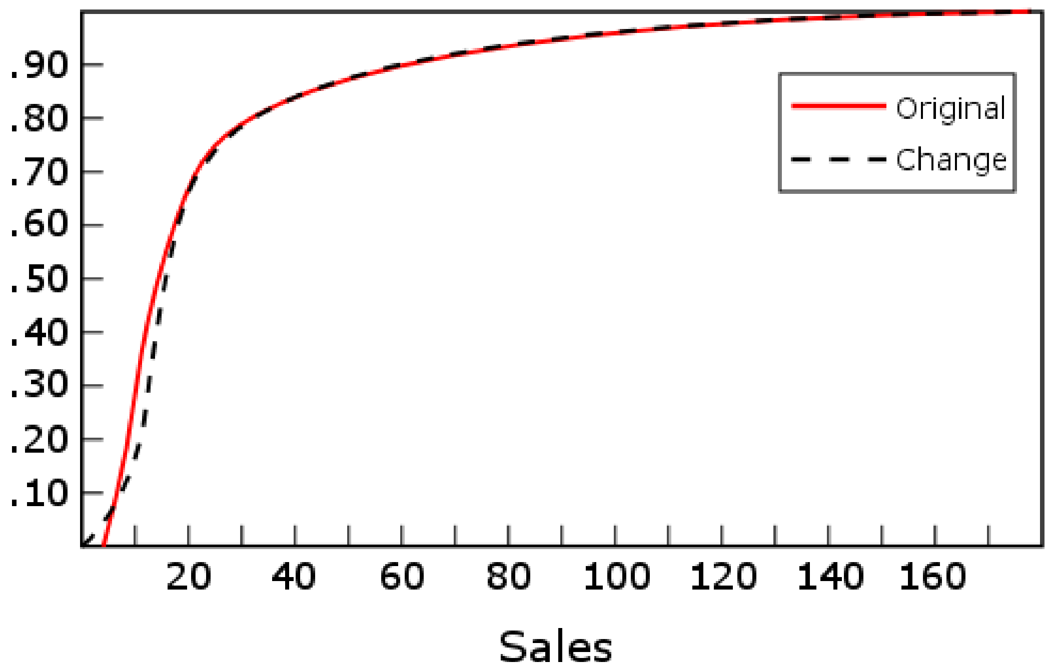

Figure 14: Change in the CDF of simulated POD sales for an increase in σ_v.

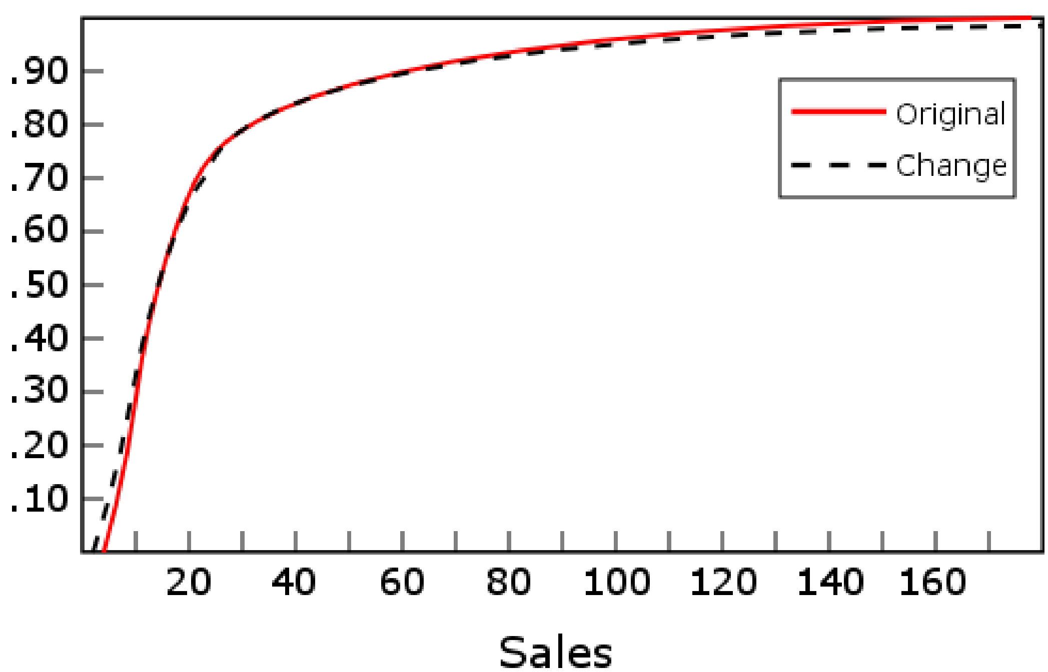

Figure 15: Change in the CDF of simulated POD sales for an increase in σ_ε.

References

Anderson, Chris. 2004. “The Long Tail.” Wired Magazine, October 2004.

Blackburn, D. (2004); “On-line Piracy and Recorded Music Sales”, Harvard University Working Paper

Brynjolfsson, E., Y. Hu and M. Smith (2003); “Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers”, Management Science, vol. 49, no. 11, pp. 1580-1596.

Brynjolfsson, Erik, Hu, Yu Jeffrey and Simester, Duncan,Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales(November 2007). Available at SSRN: http://ssrn.com/abstract=953587

Chevalier, J. and A. Goolsbee (2003); “Measuring prices and price competition online: Amazon vs. Barnes and Noble”, Quantitative Marketing and Economics, vol. 1, no. 2, pp. 203-222.

Crampes, C. and Langinier, C."Litigation and Settlement in Patent Infringement Cases" RAND *Journal of Economics*, Vol.33.(2002), pp. 258-274

Choi, J.P. "Patent litigation as an Information-Transmission Mechanism." *American Economic Review*, Vol. 88 (1998), pp.258-274

Elberse, Anita (2008). "Should You Invest in the Long Tail?" Harvard Business Review July-August 2008.

Fudenberg, D. and Tirole, J. "Game Theory" Massachusetts Institute of Technology (1991) pp.321-326

Hay, B.L. "Effort, Information, Settlement, Trial." *Journal of Legal Studies*, Vol. 24 (1995), pp. 29-62

Herald, Paul. "Optimal Remedies for Patent Infringement: A Transactional Model" *The University of Chicago* (Working Paper)

Huberman, Loch, and Onculer (2004). "Status as a valued resource", Social Psychology Quarterly, vol. 67, no. 1, pp. 103-114.

Huberson, Bernardo, Daniel Romero and Fang Wu (2008). Crowdsourcing Attention and Productivity. Working Paper.

Kreps, D. and Wilson, R. "Reputation and Imperfect information" *Journal of Economic Theory.* Vol. 27 (1982), pp. 253-279

Lanjouw, J. and Lerner, J "Preliminary Injunctive relief: Theory and Evidence from Patent Litigation" *Harvard Business School* (Working Paper).

Lanjouw, J.O. and Schankerman, M. "Characteristics of patent litigation: a window on competition" *Rand Journal of Economics*, Vol. 32 (2001), pp.129-151

Lerner, J. "Patenting in the Shadow of Competitors" *Journal of Law and Economics*, Vol. 38 (1995), pp. 463-496

Mailath, G.J. and Samuelson, L. "Who Wants a Good Reputation?" *Review of Economic Studies* Vol. 68 (2001) pp.415-441

Meurer, M.J. "The Settlement of Patent Litigation." *RAND Journal of Economics*, Vol. 20 (1989), pp. 77-91

Osborne, M.J. and Rubinstein, A. "A Course in Game Theory" The MIT Press (1994).

Pakes, A. (1986); "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks", Econometrica, vol. 54, no. 4, pp. 755-784

Pakes, A. and D. Pollard (1989); "Simulation and the Asymptotics of Optimization Estimators", Econometrica, vol. 57, no. 5, pp. 1027-1057.

Pincus, Laura B. "The Computation of Damages in Patent Infringement Actions" *Harvard Journal of Law and Technology* Vol.5 (1991), pp. 95-143

Reich-Hale, David, "Publisher Tests Advance Prints with Amy Fisher Book" Long Island Business News, September 24, 2004.

Rosenthal, M. (2004); Print-on-Demand Book Publishing, Foner Books.

United States Court of Appeals. "Alcatel-Lucent v. Microsoft" (2008)