Strategic Search in Patenting

Nick Dunagan*

July 23, 2026

Abstract

I develop a three-player game-theoretic model of patenting. An *applicant* chooses how extensively to search for prior art, determining the information brought into the patenting process and the claimed patent breadth. An *examiner* conducts additional search and might narrow the claimed breadth. A *challenger* conducts additional search and might challenge the patent. The public record from earlier search affects the likelihood of challenge by affecting the informativeness of later search. My equilibrium characterization yields comparative statics for how search and breadth respond to challenger filing costs, procedural standards, and examiner capacity. These results generate testable implications for patent breadth, grant outcomes, and post-grant challenge behavior.

Keywords: patent breadth; strategic search; information acquisition; patent prosecution; examiner search; post-grant challenge; patent quality

JEL Codes: O34, D83, K11

*John E. Walker Department of Economics, Clemson University, SC 29634 (email: nddunag@g.clemson.edu)

1 Introduction

The U.S. patent system is meant to reward inventive activity while limiting patent rights to the inventor's actual contribution. Overly narrow rights may provide too little incentive to innovate, while overly broad rights can distort competition, deter follow-on invention, and create uncertainty for rivals operating near existing technologies. The economic problem is therefore whether the patent system produces rights with appropriate breadth.

A patent's breadth is determined by searches for prior art by the applicant, the patent examiner, and later challengers. Before filing, applicants may search for prior art and decide how aggressively to pursue broad claims. After filing, examiners conduct additional search, evaluate novelty and non-obviousness, and may narrow the application. After grant, competitors or other interested parties may invest in additional search and decide whether to challenge the patent in court or through administrative review.

Existing research analyzes different parts of the strategic search problem in patenting, but typically studies them in isolation. For example, Atal and Bar show applicants may have little incentive to search when search can reduce the probability of obtaining a patent, though applicant search can rise with examiner effort, patenting fees, and R&D costs [2]. Langinier and Marcoul show applicant revelation and examiner scrutiny are jointly determined, with examiner screening intensity responding to the information transmitted by the applicant [21]. Their later work shows monetary incentives and career concerns can intensify examiner search effort, underscoring that search incentives are endogenous and policy-sensitive on both sides of the patenting process [20].

I develop a three-player model providing a unified account of how strategic search by an applicant, examiner, and challenger jointly determines patent breadth. Applicant search determines how much information enters prosecution and how broad the application can plausibly be. Examiner search determines how much claimed breadth survives into the granted patent. The prosecution record then shapes future challengers' beliefs about whether additional search and filing are worthwhile.

Focusing on equilibrium breadth rather than validity is important because a patent's economic effects often depend on how far its claims extend beyond the invention's legally defensible core. A patent may be granted, licensed, asserted, and influence rivals' behavior long before validity is fully tested. Empirical work on claim structure and examination supports this emphasis: examination tends to narrow claims, broader initial claims are associated with different grant and prosecution outcomes, and examiner-driven variation in patent rights has meaningful downstream effects on value, litigation, and innovation [28, 10].

My model yields comparative statics characterizing how applicant search, examiner scrutiny, challenger search, filing behavior, and equilibrium breadth respond to changes in institutional features. For example, a change in the cost of challenging patents affects the expected payoff from seeking aggressive breadth in the first place. If challenges become cheaper or more credible, applicants may have less incentive to pursue claims likely to be narrowed later and more incentive to search earlier to understand the limits of valid protection. Examiners may also respond differently when the downstream environment changes. The comparative statics also speak to examiner capacity, search requirements, and others policies that alter incentives to acquire information at different stages of the patenting process.

A natural application of the model is the reduction in challenger filing costs associated with Patent Trial and Appeal Board (PTAB) review. The America Invents Act (AIA) created new administrative routes for challenging issued patents, including post-grant review and inter partes review. These changes to the patenting process, alongside other policy shifts, constitute situations for empirical investigation.

Recent research suggests changes to downstream challenge institutions can shape applicant and examiner incentives before rights are granted. Schankerman and Schuett study screening for weak patents in a theoretical framework with endogenous challenge and show courts, even if perfect, cannot fully solve the screening problem because many weak patents are never challenged; they also show PTAB's welfare effect is theoretically ambiguous once indirect strategic effects are included [33]. Empirical evidence from Martinelli and Mazzei suggests post-grant review can generate upstream discipline: after the AIA, applicants in affected settings narrow patent breadth more aggressively, and examiners become tougher when applicants do not do so themselves [29]. Related empirical work shows examiner time constraints affect scrutiny and grant behavior, and greater ex ante examination effort can improve the validity of issued patents [12, 11].

Overall, patent breadth is best understood as the equilibrium outcome of a multi-stage strategic search process. Applicants do not choose claim breadth in isolation, examiners do not screen applications in a vacuum, and challengers do not evaluate patents independently of the information generated during prosecution. Each stage shapes the next. By bringing these choices together, the model provides a framework for studying how pre-grant and post-grant institutions jointly determine the breadth of patent rights.

The paper proceeds as follows. Section 2 describes the U.S. patenting process while Section 3 situates the paper within the patenting literature. Sections 4 and 5 respectively develop the model and characterize equilibrium behavior. Section 6 derives comparative statics with respect to challenge costs and other policy parameters. Section 7 discusses empirical implications, with particular attention to post-grant review and patent breadth. Section 8 concludes.

2 The U.S. Patenting Process

A patent application begins with a specification and claims describing the invention and defining the protection sought. The application must adequately disclose the invention, and later amendments must remain supported by the original filing; new matter cannot be added after filing.

The application enters *prosecution* after it's filed. Applicants may submit prior-art references through an information disclosure statement, respond to examiner objections with arguments and amendments, provide supporting evidence, and participate in examiner interviews that become part of the record. The Patent Office evaluates the claims against the relevant legal standards, then it decides whether the proposed breadth should be allowed, narrowed, or rejected [43, 39, 42].

Applicants and their representatives owe a duty of candor and good faith that requires them to disclose known information material to patentability, but they have no general duty to conduct exhaustive prior-art searches. That is, applicants must reveal what they have learned about prior art, but they are not required to learn anything. Consequently, the applicant-generated information entering prosecution is endogenous [40, 39].

This distinction between disclosure and search matters for patent breadth. The breadth an applicant can plausibly pursue depends on the prior art already known to the applicant and placed into the prosecution record. When applicants search more extensively, they may uncover constraining prior art before examination begins. When they search less, the examiner may inherit a thinner record and broader proposed claims. Applicant search therefore affects both the information available during prosecution and the breadth carried into examination.

Examiner search is the next institutional margin shaping patent breadth. After reviewing the specification and claims, the examiner searches the prior art and evaluates patentability in light of the record. The examiner must consider prior art cited by the applicant or uncovered independently, and it should reject claims the prior art and record show are unlikely to be patentable. Examiner search and review therefore help determine how much of the breadth carried into prosecution survives into an allowed patent [41].

Finally, third parties may challenge issued claims through litigation or through administrative proceedings before the Patent Trial and Appeal Board. Post-grant review begins with a third-party petition within nine months of grant and may raise a broad set of invalidity grounds. Inter partes review allows review of one or more claims on prior-art grounds under Sections 102 or 103. These procedures make post-grant search and challenge part of the broader institutional environment in which patent breadth is determined [38].

The key institutional feature connecting these stages is sequential information production that shapes patent breadth. Applicants generate and disclose some information before and during prosecution. Examiners add their own search and evaluate the claims in light of the combined record. Potential challengers observe the issued patent and prosecution record before deciding whether additional search and challenge are worthwhile. Patent breadth is therefore shaped by information generated at multiple stages, with each information acquisition decision being strategic.

3 Literature Review

As this paper treats patent breadth as the central outcome of the patenting process, it first relates to the literature on the economic consequences of patent breadth. Classic work studies the tradeoff between patent breadth and innovation incentives, showing that broader rights can increase rewards to inventors while also imposing distortions on rivals and follow-on innovators [19, 35, 13]. Empirical work likewise connects variation in patent breadth and informational content to private value. Broader patents are associated with higher firm valuations, patent citations are informative about market value, and procedural indicators such as family size and opposition outcomes are related to the value of patent rights [24, 14, 15]. More recent work moves closer to the object studied here by measuring patent breadth from claims data and showing that examination tends to narrow claims, while examiner-driven variation in patent rights affects value, litigation, and innovation [28, 10]. This literature establishes both why patent breadth matters economically and why it is useful to treat breadth as an outcome of the prosecution process rather than as a fixed feature of an invention.

The prosecution record that shapes breadth begins with applicant-side information production.

Atal and Bar show that applicants may have weak incentives to search for prior art when search can reduce the probability of obtaining a patent, although those incentives rise with examiner effort, patenting fees, and R&D costs [2]. Langinier and Marcoul extend this logic by showing that applicant disclosure and examiner scrutiny can be jointly determined, with examiner behavior responding to the information applicants reveal [21]. Empirical work on citations in prosecution reinforces the importance of distinguishing applicant-generated information from the total body of prior art. Examiners add a substantial share of citations in issued patents, applicant-submitted citations are often not central to actual rejections, and examiner citations appear more predictive of patent value than applicant citations [1, 7, 16]. Legal scholarship on inequitable conduct further highlights the distinction between acquiring information and withholding information already known, while doctrinal changes in *Exergen* and *Therasense* altered the conditions under which withholding can generate liability [6, 8, 36]. Together, this work motivates treating the applicant-generated record as a distinct endogenous input into prosecution.

The prosecution record is then augmented through examiner search and screening. Early empirical work shows that examiners are heterogeneous and that examination outcomes vary systematically with examiner characteristics rather than reflecting a uniform screening technology [5, 22]. Theoretical work frames patent examination as an incentive problem in which examiners must exert effort and truthfully reveal what they find, implying that compensation, auditing, and career concerns can affect screening quality [34, 20]. More recent empirical work shows that institutional constraints matter in practice. Reducing available review time lowers scrutiny and increases grants on the margin, greater examination resources can improve validity in high-stakes settings, and heavier workloads are associated with grant bias and lower examination quality [12, 11, 18]. This literature establishes that examiner screening is endogenous and policy-sensitive, while also highlighting the importance of examination for the breadth and quality of issued patent rights.

The final stage of the process concerns weak patents, downstream challenge, and post-grant review. Lemley and Shapiro characterize patents as probabilistic rights rather than determinate exclusion rights, while Farrell and Shapiro show how uncertainty about patent validity affects licensing and welfare [23, 9]. Schankerman and Schuett bring examination and challenge together in a model of patent screening with endogenous court challenges, showing that downstream courts cannot fully solve the screening problem because many weak patents are never challenged [33]. On the institutional side, the America Invents Act created new administrative challenge routes, especially inter partes review and post-grant review, lowering the cost of contesting some issued patents. Early evidence documents the practical importance of IPR, while later work uses PTAB outcomes to study the determinants of patent quality more directly [26, 27]. Most closely related to the present paper, Martinelli and Mazzei show that post-grant review can generate upstream effects: after the AIA, applicants narrowed patent breadth more aggressively and examiners became tougher when applicants did not do so [29]. This literature demonstrates that post-grant institutions can affect incentives before rights are granted.

Taken together, these literatures show that applicant-side information production, examiner screening, and downstream challenge each affect the breadth and quality of patent rights. Applicant-side work explains why incentives to search for and reveal prior art may be weak or distorted [2, 21]. Examiner-side work shows that scrutiny responds to office incentives, resources, and institutional design [34, 12, 11]. Challenge-side work explains why downstream review is selective and why post-grant institutions can affect behavior during prosecution [33, 29]. This paper builds on these insights by studying patent breadth as the outcome of strategic search across the patenting process. Rather than treating pre-grant search, examination, and post-grant review as separate margins, the framework emphasizes how incentives to acquire information at one stage can shape behavior at later stages, and how changes in post-grant institutions can feed back into claim breadth and examination before grant.

4 The Model Structure

This section presents a finite-horizon, three-player dynamic game among an *applicant*, an *examiner*, and a potential *challenger*. The central object of interest is patent breadth, which is shaped by players' costly information acquisition before and after grant.

The game has seven steps grouped into the *applicant stage*, the *examiner stage*, and the *challenger stage*. The applicant begins with initial information and chooses whether to search further; the resulting set of knowledge determines the breadth carried into examination. The examiner observes the submitted record, chooses whether to search further, and truncates breadth according to the information in the record. The applicant then decides whether to accept or abandon the resulting patent. The potential challenger observes the granted breadth and prosecution record, chooses whether to search further, then decides whether to file a challenge. Challenge success is determined by the patent's true overbreadth.

4.1 Environment, Timing, and Payoffs

Several parameters in the game are common knowledge. The patent's quality $q \in Q \subset \mathbb{R}_+$ acts as a scaling object on the value of patent breadth to the applicant and challenger. This formulation differs from those in which quality is an endogenous valuation object [31, 25, 32]. The parameters λ_A, λ_E, and λ_C, affect each player's marginal cost of search.

Other parameters are not common knowledge. No player observes the patent's legally valid breadth $b^* \in B \subset \mathbb{R}_+$ which Nature draws from a commonly known distribution F_0. The applicant and challenger each privately observe their respective legal costs$K_A \sim G_A$ and $K_C \sim G_C$, which are incurred only if a challenge occurs after grant. Finally, the applicant privately learns its initial information X_0, which represents research conducted before the beginning of the game.

Search arises in the model because the legally valid breadth remains uncertain at each player's informational starting point. Although the initial information set X_0 provides partial information about b^*, it does not reveal b^* itself. This uncertainty matters because patent value depends on the breadth ultimately secured, while the value and likely success of a later challenge depend on the relationship between granted breadth and legally valid breadth. The applicant therefore chooses search from the informational position defined by X_0, and the record carried into examination consists of this initial information together with any additional information uncovered through applicant search.

The timing of the game is as follows.

1. Nature draws b^*, K_C, K_A, and X_0.

2. The applicant chooses search effort s_A, which generates post-search information summarized by X_A. The claimed breadth is the minimum of X_A.

3. The examiner observes X_A and chooses search effort s_E, which updates the information state to X_E. The granted breadth b_E is the minimum of X_E.

4. The applicant chooses whether to accept the granted breadth or abandon the patenting process.

5. Conditional on grant, the challenger observes the granted breadth and the prosecution record, then chooses search effort s_C.

6. After observing the outcome of their search, the challenger decides whether to file a challenge.

7. If a challenge is filed, the true breadth b^* is revealed and challenge success is probabilistically determined by the realized overbreadth $b_E - b^*$.

My model abstracts from the applicant's choice of which discovered references to place into the prosecution record to focus on how earlier search shapes downstream challenge incentives. [1] This seems a reasonable starting point, given the disclosure duties and possible penalties for withholding material information.

Ignoring strategic disclosure is especially natural when the expected penalty from concealing known material information is large. In patent law, the relevant deterrent is inequitable conduct: if an applicant knowingly withholds material prior art with deceptive intent, the resulting patent may be rendered unenforceable, and the attorneys involved may face professional or legal consequences. In that environment, the applicant's most relevant strategic margin is not whether to disclose discovered prior art, but whether to acquire that information in the first place. Once a material reference is known, withholding it may create legal risk large enough to dominate the private benefit of concealment. The deterrence problem therefore operates primarily through search: by choosing how much prior art to look for, the applicant affects both the breadth initially supported by the record and the information carried into prosecution. The model captures this search-as-deterrence channel while leaving a richer model of selective disclosure and concealment penalties for future work.

The players differ in what they value from the patenting process. If a patent is granted at breadth b and survives any post-grant challenge, the applicant receives value

$$V_A(b, q) = v(q)u(b),$$

where $v'(q) > 0, u'(b) > 0$, and $u''(b) < 0$. Patent quality scales the value of the right, while breadth has positive but diminishing marginal value. The applicant's return is evaluated net of any search costs and legal costs it incurs.

The challenger's value from invalidating an overly broad patent is

$$V_C(b, q) = v(q)u(b),$$

[1] Including a disclosure stage would make the model richer, but it would also combine the search mechanism with a separate concealment problem, significantly detracting from the tractability of the model.

using the same basic value structure as the applicant. This symmetry is not essential, but it captures the idea that broader and higher-quality patents are more valuable both to hold and to remove as competitive constraints. The challenger's return is evaluated net of any search costs and legal costs it incurs.

The examiner does not value patent breadth directly, but dislikes granting claims exceeding the legally valid breadth. The examiner's loss is

$$V_E(b_E - b^*),$$

where $V_E(0) = 0$, $V_E'(b_E - b^*) > 0$, and $V_E''(b_E - b^*) > 0$. Thus, the examiner's return is evaluated net of any search costs it incurs.

The probability a challenge succeeds is

$$p_C(b^*, b_E) = \rho(b_E - b^*),$$

where $\rho : \mathbb{R}_+ \to [0, 1]$ is weakly increasing in realized overbreadth.

There are four outcomes to identify payoffs. The examiner's payoff does not change based on outcome, but always has a payoff of combined overbreadth loss and search costs. First, if b_E is declined, then the applicant incurs search costs with no benefit. Second, if b_E is accepted and not challenged, then the applicant receives their value of the patent and incurs search costs while the challenger incurs any search costs with no benefit. Third, if b_E is accepted and challenged successfully, then the applicant receives no benefit but incurs search and legal costs while the challenger receives their value of the patent net of any search and legal costs. Finally, if b_E is accepted and challenged unsuccessfully, then the applicant receives their value of the patent net of search and legal costs while the challenger receives no benefit but incurs any search and legal costs.

4.2 Beliefs and the Truncation-Draw Information Structure

Search generates observations that each can be interpreted as a piece of prior art placing an upper bound on the patent's legally valid breadth b^*. Formally, conditional on b^*, each search process generates a finite set of constraining draws

$$X_t = \{x_1, \ldots, x_s\}, \qquad x_i \geq b^*., \qquad \text{for} \quad t \in \{0, A, E, C\}$$

For any realized record X_t, define

$$m(X_t) = \min(X_t), \quad \text{and} \quad R(X_t) = |X_t|, \qquad \text{for} \quad t \in \{0, A, E, C\}$$

I write these objects simply as m_t and R_t. The statistic m is the tightest observed upper bound on b^*, while R is the size of the accumulated record. The model assumes players condition on the record only through these two statistics, so beliefs about b^* are summarized by the probability distribution

$$F(\cdot \mid m, R),$$

with posterior mean

$$\mu(m, R) \equiv E[b^* \mid m, R].$$

The two statistics affect beliefs through distinct channels. The minimum m creates a boundary effect: because each constraining draw provides an upper bound on b^*, posterior support is contained in $[0, m]$, and a lower minimum shifts beliefs toward lower legally valid breadth. For any fixed granted breadth, this shift tends to increase expected overbreadth. Record size R instead creates a screening effect: holding m fixed, a larger record makes the observed minimum more informative and shifts beliefs toward greater concentration near m, tending to reduce expected overbreadth. The following assumptions formalize these posterior properties.

Assumption 1 ***Support restriction:*** *For every admissible* (m, R)*,*

$$\operatorname{supp}\bigl(F(\cdot \mid m, R)\bigr) \subseteq [0, m].$$

Assumption 2 ***Minimum monotonicity:*** *Holding* R *fixed, a lower minimum* m *shifts posterior beliefs toward lower values of the true legally valid breadth* b^* *in the sense of first-order stochastic* *dominance. Formally, for any $m' < m$,*

$$F(b \mid m', R) \geq F(b \mid m, R) \qquad \textit{for all } b \in B.$$

Assumption 3 ***Record-size monotonicity:*** *Holding m fixed, a larger record shifts posterior beliefs toward greater concentration near m. Formally, if $R' > R$, then*

$$\mu(m, R') \geq \mu(m, R), \quad \textit{with} \quad \lim_{R \to \infty} \mu(m, R) = m.$$

Much of the model's comparative-static logic follows from the interaction between these boundary and screening effects. Figure 1 illustrates these two belief channels.

The belief structure is similar at each stage, but with different realized records. The applicant begins from (m_0, R_0), their search produces (m_A, R_A), and the breadth entering examination is m_A. Examiner search updates the record to (m_E, R_E), and the granted breadth $b_E = m_E$. Conditional on grant, challenger search updates the prosecution record to (m_C, R_C). At every stage, posterior beliefs over b^* are summarized by the current pair (m, R).

The truncation-draw structure reflects the stochastic nature of prior-art discovery. Neither research and development nor prior-art search proceeds deterministically toward a known legal boundary. When a player searches, the player does not know whether the search will uncover a new constraint that materially changes the patent's legally supportable breadth. A deterministic search technology would mechanically translate effort into narrower breadth or more precise beliefs, while a direct-revelation technology would make search a probabilistic discovery of b^*. The truncation-draw structure instead captures the institutional setting in which search uncovers pieces of prior art that impose upper bounds on valid breadth, allowing search either to lower the most constraining observed reference or to expand the record without changing that minimum.

Search is discrete for both substantive and technical reasons. Substantively, prior-art search is often organized in discrete increments such as search queries, attorney hours, database searches, examiner time allocations, or challenger investigation steps. Technically, a finite search space ensures each player's optimization problem admits a solution without imposing stronger smoothness, concavity, or single-crossing assumptions.

Because a marginal increase in search can leave the minimum unchanged while enlarging the

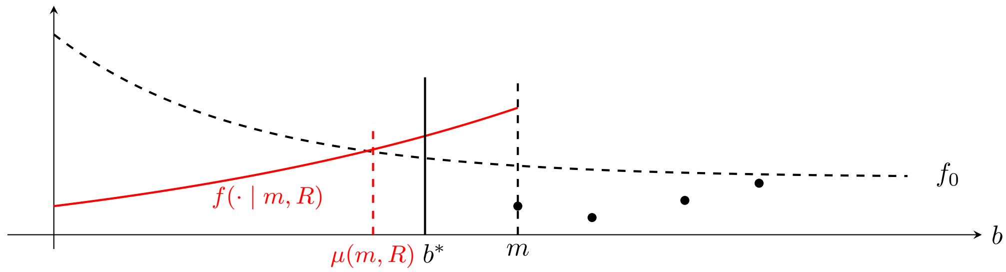

(a) Small record

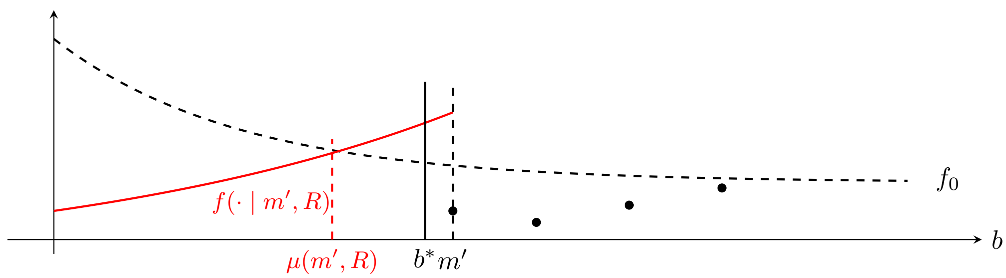

(b) Same record size, smaller minimum

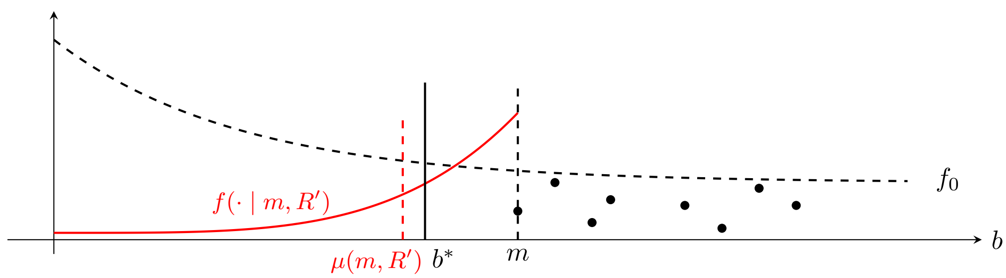

(c) Same minimum, larger record

Figure 1: Truncation-draw belief updating. Each realized draw is a constraining observation and therefore lies weakly above the true legally valid breadth b^*. Panel (a) shows a baseline record with minimum draw m and posterior density $f(\cdot \mid m, R)$. Panel (b) keeps record size fixed but lowers the minimum to $m' < m$, which lowers the upper bound on feasible legally valid breadth and shifts the posterior mean downward. Panel (c) holds the minimum fixed at m but increases the number of observed draws from R to $R' > R$. The larger record makes the observed minimum more informative and shifts posterior beliefs toward m.

record, or lower it discontinuously, continuation values need not be differentiable and best responses may be set-valued. The equilibrium analysis therefore uses monotone comparative statics rather than first-order conditions. The Topkis-style arguments below characterize monotone selections from these best-response correspondences without requiring differentiability of the value functions [37, 30].

5 Equilibrium Characterization

I solve the game via backward induction, beginning with the challenger stage. I first characterize the challenger's filing and search decisions, then derive the applicant's accept-or-abandon rule, the examiner's search decision, and the applicant's initial search decision. The resulting strategies form a subgame-perfect Nash equilibrium, which is suitable although the applicant's defense cost K_A and the challenger's filing cost K_C are privately observed. This is because they do not enter the payoffs of other players directly. They affect continuation play only through the actions they induce, which earlier-moving players evaluate using the known cost distributions.

5.1 The Challenger Stage

Conditional on the applicant accepting the granted patent, the challenger observes the inherited prosecution record (m_E, R_E) and the granted breadth $b_E \equiv m_E$. The challenger chooses whether to acquire additional information through search, then observes the record (m_C, R_C), and decides whether to file a challenge. Consequently, I characterize the filing rule first to define the value of search.

5.1.1 Challenger Filing

For a realized post-search information state (m_C, R_C), the challenger's posterior-expected probability of success is

$$\bar{p}_C(b_E; m_C, R_C) \equiv \mathbb{E}[\rho(b_E - b^*) \mid m_C, R_C] \, .$$

The challenger files if and only if the expected payoff from filing is non-negative, yielding the threshold rule

$$K_C \leq \hat{K}_C(b_E; m_C, R_C) \equiv V_C(b_E, q)\bar{p}_C(b_E; m_C, R_C).$$

Thus, conditional on the post-search state, filing is characterized by a cutoff in the challenger's privately observed challenge cost. If a challenge is filed, its realized probability of success is $\rho(b_E - b^*)$.

5.1.2 Challenger Search

Before observing the outcome of additional search, the challenger chooses search intensity

$$s_C \in S = \{0, 1, \ldots, \bar{s}\},$$

with cost $c_C(s_C; \lambda_C)$, where search costs are increasing and convex in effort. Search generates additional constraining draws

$$X_C(s_C) = \{x_{C1}, \ldots, x_{C,s_C}\},$$

which are combined with the inherited prosecution record to produce

$$m_C(s_C) = \min\bigl(\{m_E\} \cup X_C(s_C)\bigr), \qquad R_C(s_C) = R_E + s_C.$$

The challenger's search problem is therefore choosing s_C to maximize

$$U_C(s_C; \theta_2) = \mathbb{E}\bigl[\max\bigl\{0, V_C(b_E, q)\bar{p}_C\bigl(b_E; m_C(s_C), R_C(s_C)\bigr) - K_C\bigr\} \bigm| m_E, R_E\bigr] - c_C(s_C; \lambda_C),$$

where

$$\theta_2 = (q, b_E, m_E, R_E, \lambda_C, K_C)$$

denotes the challenger's state. Search is valuable because it changes the distribution of post-search states (m_C, R_C), and therefore the expected option value of filing.

5.1.3 Existence of a Challenger Search Best Response

The challenger's search problem admits a solution because search is chosen from the finite set

$$S = \{0, 1, \ldots, \bar{s}\}.$$

Under the maintained assumptions on the search technology and search costs, $U_C(s_C; \theta_2)$ is well defined for every $s_C \in S$. Because the challenger maximizes a real-valued function over a finite set, the best-response correspondence

$$BR_C(\theta_2) \equiv \arg\max_{s_C \in S} U_C(s_C; \theta_2)$$

is nonempty. The best-response correspondence might have multiple elements, so I work with the monotone selections

$$s_C^L(\theta_2) \equiv \min BR_C(\theta_2) \quad \text{and} \quad s_C^H(\theta_2) \equiv \max BR_C(\theta_2).$$

When I refer to equilibrium challenger search below, s_C^* denotes one such selection.

5.2 The Applicant Accept or Abandon Decision

The applicant's accept-or-abandon decision anticipates the challenger stage. After examiner search, the applicant observes the prosecution record (m_E, R_E) and granted breadth $b_E \equiv m_E$, then decides whether to accept the patent or abandon the process.

Two probabilities summarize the downstream challenge risk. First, because K_C is privately observed by the challenger but drawn from the known distribution G_C, earlier players perceive challenge probabilistically. Given equilibrium challenger search s_C^*, the ex ante probability of filing is

$$p_f(b_E; m_E, R_E) = \mathbb{E}_{X_C | s_C^*} \left[G_C\Big(\hat{K}_C(b_E; m_C, R_C)\Big) \Big| m_E, R_E \right].$$

Second, the applicant's perceived probability of surviving on the merits if challenged is

$$\bar{p}_A^M(b_E; m_E, R_E) \equiv \mathbb{E}[1 - \rho(b_E - b^*) \mid m_E, R_E].$$

Let $V_A \equiv V_A(b_E, q)$. The applicant's expected payoff from accepting the patent is

$$(1 - p_f)V_A + p_f \left(\bar{p}_A^M V_A - K_A\right).$$

The applicant accepts if and only if this expected payoff is non-negative. For $p_f > 0$, rearranging yields the threshold rule

$$\hat{K}_A \equiv V_A \left(\frac{1}{p_f} + \bar{p}_A^M - 1 \right) \geq K_A.$$

Thus, the applicant accepts when the value of the granted patent, adjusted for challenge risk and expected merits survival, is sufficiently large relative to defense cost. If $p_f = 0$, the applicant accepts whenever $V_A \geq 0$.

5.3 The Examiner Stage

The examiner moves after applicant search and before the applicant's accept-or-abandon decision. After observing the applicant-generated record (m_A, R_A), the examiner chooses additional search. Search generates s_E additional constraining draws X_E, which are combined with the inherited record to produce

$$m_E(s_E) = \min(\{m_A\} \cup X_E), \qquad R_E(s_E) = R_A + s_E.$$

Under the truncation rule, granted breadth is determined by the most constraining information in the combined record:

$$b_E(s_E) = m_E(s_E).$$

5.3.1 Examiner Search

Given the inherited state (m_A, R_A), the examiner chooses search intensity

$$s_E \in S = \{0, 1, \ldots, \bar{s}\},$$

with cost $c_E(s_E; \lambda_E)$, where search costs are increasing and convex in effort. The examiner's objective is to minimize expected overbreadth loss net of search costs:

$$U_E(s_E; \theta_1) = -\mathbb{E}[V_E(b_E(s_E) - b^*) \mid m_A, R_A] - c_E(s_E; \lambda_E),$$

where

$$\theta_1 = (m_A, R_A, \lambda_E)$$

denotes the examiner's inherited state. Thus, examiner search is valuable only insofar as the induced post-search record reduces expected overbreadth enough to justify the additional search cost.

5.3.2 Existence of an Examiner Search Best Response

The examiner's search problem, like the challenger's, involves maximizing a real-valued function over a finite set. Therefore, the best-response correspondence

$$BR_E(\theta_1) \equiv \arg\max_{s_E \in S} U_E(s_E; \theta_1)$$

is nonempty, but might contain multiple elements. I therefore work with the monotone selections

$$s_E^L(\theta_1) \equiv \min BR_E(\theta_1) \quad \text{and} \quad s_E^H(\theta_1) \equiv \max BR_E(\theta_1).$$

When I refer to equilibrium examiner search below, s_E^* denotes one such selection.

5.4 The Applicant Stage

The applicant moves first strategically. The applicant begins from the initial information state (m_0, R_0), observes patent quality q, search-cost shifter λ_A, and defense cost K_A, and chooses search intensity s_A. Recall, there is no separate disclosure or breadth-choice stage.

5.4.1 Applicant Search

The applicant chooses

$$s_A \in S = \{0, 1, \ldots, \bar{s}\},$$

with cost $c_A(s_A; \lambda_A)$, where search costs are increasing and convex in effort. Search generates additional constraining draws X_A, which are combined with the applicant's initial information to form

$$X_1 = X_0 \cup X_A.$$

The applicant-generated record entering examination is summarized by

$$m_A(s_A) = \min X_1, \qquad R_A(s_A) = |X_1|.$$

Applicant search is a value-of-information problem over the continuation game beginning from the state (m_A, R_A). Let

$$\theta_0 = (q, \lambda_A, K_A, m_0, R_0).$$

Because the applicant can abandon after examination, the continuation value from the applicant-generated state (m_A, R_A) is

$$Z(m_A, R_A; \theta_0) = \mathbb{E}\left[\max\left\{0, V_A(b_E, q)\left[1 - p_f(b_E, q; m_E, R_E)\left(1 - \bar{p}_A^M(b_E; m_E, R_E)\right)\right] - p_f(b_E, q; m_E, R_E)K_A\right\} \,\middle|\, m_A, R_A\right]$$

where the expectation is over examiner and challenger continuation play induced by (m_A, R_A).

The applicant's ex ante search problem is therefore to choose s_A to maximize

$$U_A(s_A; \theta_0) = \mathbb{E}\left[Z\left(m_A(s_A), R_A(s_A); \theta_0\right) \,\middle|\, m_0, R_0\right] - c_A(s_A; \lambda_A),$$

where the expectation is over the random outcome of applicant search. Thus, applicant search is valuable because it changes the continuation value of the prosecution and challenge game. Unlike examiner search, which directly reduces expected overbreadth, or challenger search, which changes the expected value of filing, applicant search shapes the information state inherited by all later players.

5.4.2 Existence of an Applicant Search Best Response

The applicant's search problem, like the challenger;s, involves maximizing a real-valued function over a finite set. Therefore, the best-response correspondence

$$BR_A(\theta_0) \equiv \arg\max_{s_A \in S} U_A(s_A; \theta_0)$$

is nonempty, but might contain multiple elements. I therefore work with the monotone selections

$$s_A^L(\theta_0) \equiv \min BR_A(\theta_0) \quad \text{and} \quad s_A^H(\theta_0) \equiv \max BR_A(\theta_0).$$

When I refer to equilibrium applicant search below, s_A^* denotes one such selection.

5.5 Equilibrium

The stagewise characterizations above can be combined to construct a subgame-perfect Nash equilibrium.

Proposition 1 *Suppose the search-action set S is finite and the search technology, cost functions, and posterior system $F(\cdot \mid m, R)$ are such that continuation payoffs are well-defined at every admissible history. There exists an equilibrium in which applicant, examiner, and challenger search are chosen from their respective best-response correspondences, granted breadth satisfies $b_E = m_E$, applicant acceptance and challenger filing are governed by threshold rules, and beliefs about the true legally valid breadth are given by $F(\cdot \mid m, R)$ for every information record.*

The existence argument follows directly from backward induction. Applicant acceptance and challenger filing have already been characterized by threshold rules, while the finiteness of S ensures that each player has at least one optimal search action whenever continuation payoffs are well-defined. The resulting equilibrium strategy profile is

$$(\{s_A^*, Y_A^*\}, \{s_E^*\}, \{s_C^*, Y_C^*\}) ,$$

where s_A^*, s_E^*, and s_C^* are the players' search rules, Y_A^* is the applicant's accept-or-abandon rule, and Y_C^* is the challenger's filing rule.

The argument starts with the challenger. At every inherited post-grant state, the challenger chooses search from the finite set S. Because its continuation payoff is well-defined, the challenger has at least one optimal search choice at each such state. Conditional on the resulting post-search record (m_C, R_C), the challenger then files according to the threshold rule $K_C \leq \hat{K_C}$. These choices define s_C^* and Y_C^* at every admissible history.

Taking the challenger's continuation behavior as given, the applicant's post-examination accept-or-abandon decision is governed by the threshold rule $K_A \leq \hat{K_A}$. The examiner anticipates this decision and the subsequent challenger stage when choosing search. Because examiner search is also selected from the finite set (S), at least one optimal examiner action exists at every applicant-generated record (m_A, R_A). The examiner's search choice produces the record (m_E, R_E), and granted breadth is then determined mechanically by $b_E \equiv m_E$.

Finally, the applicant anticipates the examiner's search, its own later acceptance decision, and the challenger's continuation behavior. The applicant chooses initial search from the finite set S, so at least one optimal applicant search action exists at every initial state (m_0, R_0). Together, these sequentially optimal choices form an equilibrium. At every on or off path history, beliefs about b^* are determined by the observed record and given by $F(\cdot \mid m, R)$, where (m, R) are the record's sufficient statistics.

6 Comparative Statics

Having characterized equilibrium behavior, this section studies how the players' decisions and the resulting equilibrium outcomes vary with the model's parameters. I first derive stagewise comparative statics, which vary one argument of a player's decision problem at a time while holding the remaining arguments fixed. Next, I turn to global effects that propagate through the sequence of play.

6.1 Stagewise Comparative Statics of Challenger's Filing Decision

Comparative statics of the challenger's filing decision follow from the filing threshold conditional on the realized post-search record:

$$\hat{K}_C(b_E; m_C, R_C) = V_C(b_E, q)\bar{p}_C(b_E; m_C, R_C).$$

Holding R_C fixed, a lower m_C raises expected overbreadth and increases the filing threshold. Holding m_C fixed, a larger R_C raises expected true breadth, lowers expected overbreadth, and reduces the filing threshold.

Holding the post-search record fixed, higher patent quality q raises the filing threshold by increasing the value term $V_C(b_E, q)$. A higher granted breadth b_E also raises the threshold by increasing the value of successful challenge and the expected probability of success. Because G_C is increasing in b_E, the same changes raise or lower the conditional probability of filing.

6.2 Stagewise Comparative Statics of Challenger's Search Decision

The search-cost shifter λ_C enters the objective only through the cost function $c_C(s_C; \lambda_C)$. A higher λ_C raises the marginal cost of search and lowers the return to higher search intensity. Thus, the search objective exhibits decreasing differences in (s_C, λ_C).

Patent quality q enters the objective through the value term $V_C(b_E, q)$. A higher q raises the value of successful challenge and increases the return to higher search intensity. Thus, the search objective exhibits increasing differences in (s_C, q).

The challenge cost K_C affects search through the filing option value. A lower K_C makes filing profitable in more post-search states and raises the value of acquiring information. Equivalently, the search objective exhibits decreasing differences in (s_C, K_C).

The inherited record size R_E affects search through the screening channel. Holding m_E fixed, a larger inherited record lowers expected overbreadth and reduces the value of additional search. Thus, the search objective exhibits decreasing differences in (s_C, R_E).

The inherited minimum m_E is also the granted breadth b_E. A higher m_E raises patent value through $V_C(b_E, q)$ and relaxes the tightest inherited upper bound on b^*. Under the maintained minimum-dominance assumption, the change in the minimum dominates the induced change in the posterior expectation of true breadth, so these channels raise expected overbreadth and increase the return to search. Thus, the search objective exhibits increasing differences in (s_C, m_E), or equivalently in (s_C, b_E).

Combining these results, challenger search is weakly increasing in q and $b_E \equiv m_E$, and weakly decreasing in λ_C, K_C, and R_E. By Topkis's Theorem, there exists a monotone selection $s_C^*(\theta_2)$ with these properties [37].

The filing effects combine the direct threshold effects described above with the indirect effects of parameters on challenger search. A higher λ_C lowers challenger search and therefore reduces the induced filing threshold and filing probability. A larger inherited record R_E lowers expected success directly and also reduces search, so both channels reduce filing. A higher inherited minimum $m_E \equiv b_E$ raises filing incentives directly and, through its positive effect on challenger search, raises filing. Higher quality q raises filing directly through patent value and indirectly through search. Finally, a lower challenge-cost environment increases filing directly by lowering the cost of challenge and indirectly by increasing the value of challenger search.

Table 1 summarizes the stagewise comparative statics for challenger search, the filing threshold, and the implied filing probability. The signs in the filing-threshold and filing-probability columns incorporate both the direct effects on the threshold and any induced effects through equilibrium challenger search.

Table 1: Stagewise comparative statics for the challenger

Parameter	Equilibrium search s_C^*	Stagewise filing threshold $\hat{K}_C$	Implied filing probability p_f	Main channel
λ_C	$\downarrow$	$\downarrow$	$\downarrow$	Higher challenger search cost lowers equilibrium search, reducing the induced post-search filing value.
q	$\uparrow$	$\uparrow$	$\uparrow$	Higher patent quality raises the value of successful challenge and the value of search.
R_E	$\downarrow$	$\downarrow$	$\downarrow$	A larger inherited record lowers expected overbreadth and reduces challenger search.
$m_E \equiv b_E$	$\uparrow$	$\uparrow$	$\uparrow$	A higher inherited minimum raises patent value and increases the value of challenger search.
K_C	$\downarrow$	$\downarrow$	–	Higher realized challenge cost lowers the filing option value and reduces challenger search.
Lower challenge-cost environment	$\uparrow$	$\uparrow$	$\uparrow$	Lower challenge costs expand the set of states in which challenge is profitable and raise the value of search.
m_C	–	$\downarrow$ when $m_C \uparrow$	$\downarrow$ when $m_C \uparrow$	A lower post-search minimum raises expected overbreadth and increases expected challenge success.
R_C	–	$\downarrow$ when $R_C \uparrow$	$\downarrow$ when $R_C \uparrow$	A larger post-search record raises expected true breadth holding the minimum fixed and lowers expected overbreadth.

Notes: The filing threshold is

$$\hat{K}_C(b_E; m_C, R_C) = V_C(b_E, q)\, \bar{p}_C(b_E; m_C, R_C).$$

The signs for $\hat{K}_C$ and p_f incorporate both direct effects on the filing threshold and, for parameters that affect challenger search, indirect effects through the induced distribution of post-search states (m_C, R_C). The final column previews the induced effect on the ex ante filing probability p_f, which is defined formally in the applicant acceptance section. Since G_C is increasing, a higher filing threshold raises filing probability, while a lower filing threshold reduces it. For the realized private cost K_C, the effect on p_f is not reported because p_f integrates over the distribution G_C. For m_C and R_C, the signs refer to the conditional filing probability inside the expectation defining p_f.

6.3 Stagewise Comparative Statics of Applicant Acceptance

The applicant's accept decision depends on applicant value $V_A(b_E, q)$, the probability of challenge p_f, and the applicant's perceived probability of surviving on the merits $\bar{p}_A^M$. The relevant comparative statics for p_f follow from the challenger-stage results in Table 1. Because G_C is increasing, changes that raise the filing threshold raise p_f, while changes that lower the filing threshold reduce it. A higher challenge-cost environment reduces p_f directly by shifting the challenger cost distribution upward.

A higher challenger search-cost shifter λ_C lowers p_f while leaving $V_A(b_E, q)$ and $\bar{p}_A^M$ unchanged. Since a lower challenge probability raises $\hat{K}_A$, higher λ_C makes acceptance more likely. The same logic applies to a higher challenge-cost environment: when challenge becomes more costly on average, filing becomes less likely and the applicant's expected payoff from accepting rises.

Patent quality q affects acceptance through two channels. Higher quality raises applicant value directly, but it also raises challenger incentives and therefore increases p_f. Under the maintained quality-value dominance condition, the direct value effect dominates the induced increase in challenge risk, so higher quality raises the accept threshold. [2]

Granted breadth $m_E \equiv b_E$ also has competing effects. A broader granted patent raises applicant value directly, but it also increases challenge incentives and lowers the applicant's perceived probability of surviving on the merits. Under the maintained breadth-value dominance condition, the direct value effect dominates the induced increase in challenge risk and the decline in merits survival, so a higher m_E raises the accept threshold. [3]

The inherited record size R_E has a cleaner effect. Holding granted breadth fixed, a larger

---

[2] For $q' > q$, let

$$V_A' = V_A(b_E, q'), \qquad p_f' = p_f(b_E; m_E, R_E, q').$$

Holding (b_E, m_E, R_E) fixed, $\bar{p}_A^M$ does not change with q. The maintained quality-value dominance condition is

$$V_A' \left(\frac{1}{p_f'} + \bar{p}_A^M - 1 \right) \geq V_A \left(\frac{1}{p_f} + \bar{p}_A^M - 1 \right).$$

If this condition were violated, the induced increase in challenge risk would dominate the direct increase in patent value, and the effect of higher quality on applicant acceptance would reverse.

[3] For $m_E' > m_E$, with $b_E' = m_E'$, let

$$V_A' = V_A(b_E', q), \qquad p_f' = p_f(b_E'; m_E', R_E), \qquad (\bar{p}_A^M)' = \bar{p}_A^M(b_E'; m_E', R_E).$$

The maintained breadth-value dominance condition is

$$V_A' \left(\frac{1}{p_f'} + (\bar{p}_A^M)' - 1 \right) \geq V_A \left(\frac{1}{p_f} + \bar{p}_A^M - 1 \right).$$

If this condition were violated, the increase in challenge risk and decline in merits survival would dominate the direct increase in patent value, and broader granted breadth would lower applicant acceptance.

prosecution record raises $\bar{p}_A^M$ by shifting beliefs toward higher true breadth. It also lowers p_f through the challenger stage. Both channels raise the accept threshold and make acceptance more likely.

Finally, defense cost K_A enters directly on the right-hand side of the accept rule. A higher K_A makes acceptance less attractive and lowers the probability that the applicant accepts the patent.

Table 2 summarizes these stagewise comparative statics.

Table 2: Stagewise comparative statics for the applicant accept decision

Parameter	Accept threshold $\hat{K}_A$	Acceptance probability	Main channel
λ_C	$\uparrow$	$\uparrow$	Higher challenger search cost lowers challenger search and reduces the probability of challenge.
Higher challenge-cost environment	$\uparrow$	$\uparrow$	Higher challenge costs reduce the probability a challenge is filed.
q	$\uparrow$	$\uparrow$	Higher patent quality raises applicant value, and under quality-value dominance this direct effect dominates the induced increase in challenge risk.
$m_E \equiv b_E$	$\uparrow$	$\uparrow$	Higher granted breadth raises applicant value, and under breadth-value dominance this direct effect dominates the induced increase in challenge risk and decline in merits survival.
R_E	$\uparrow$	$\uparrow$	A larger prosecution record lowers challenge risk and raises the applicant's perceived probability of surviving on the merits.
K_A	–	$\downarrow$	Higher defense cost makes acceptance less attractive directly.

Notes: The accept threshold is

$$\hat{K}_A \equiv V_A(b_E, q)\left(\frac{1}{p_f} + \bar{p}_A^M - 1\right),$$

and the applicant accepts if and only if $\hat{K}_A \geq K_A$. The ex ante probability of challenge is p_f, and $\bar{p}_A^M$ is the applicant's perceived probability of surviving on the merits if challenged. The signs for q and $m_E \equiv b_E$ use the maintained quality-value and breadth-value dominance conditions, respectively. The defense cost K_A does not shift $\hat{K}_A$; it affects acceptance through the right-hand side of the accept rule.

6.4 Stagewise Comparative Statics of Examiner Search

I next characterize how the examiner's search choice responds to changes in the inherited record and the examiner's search cost.

The search-cost shifter λ_E enters only through the cost function $c_E(s_E; \lambda_E)$. A higher λ_E raises the cost of higher search intensity and lowers the net return to search. Thus, the search objective exhibits decreasing differences in (s_E, λ_E).

The inherited record size R_A affects the examiner's beliefs through the screening channel. Holding m_A fixed, a larger record shifts beliefs toward higher true breadth and lowers expected overbreadth. This reduces the value of additional examiner search, so the search objective exhibits decreasing differences in (s_E, R_A).

The inherited minimum m_A affects search through competing channels. A higher m_A means the patent enters examination with broader proposed breadth, but it also raises the examiner's posterior expectation of true breadth. Under the maintained minimum-dominance assumption, the increase in starting breadth dominates the induced increase in expected true breadth, so expected overbreadth rises. A higher m_A therefore increases the value of examiner search, and the search objective exhibits increasing differences in (s_E, m_A).

Combining these results, equilibrium examiner search is weakly decreasing in λ_E and R_A, and weakly increasing in m_A. By Topkis's Theorem, there exists a monotone selection $s_E^*(\theta_1)$ with these properties [37]. Because examiner search determines the post-search prosecution record and granted breadth, these responses shape the state inherited by the applicant and challenger. [4]

Table 3 summarizes the stagewise comparative statics for examiner search.

Table 3: Stagewise comparative statics for the examiner

Parameter	Equilibrium search s_E^*	Main channel
λ_E	$\downarrow$	Higher examiner search costs lower the net return to search.
R_A	$\downarrow$	A larger inherited record lowers expected overbreadth through the screening channel, reducing the value of additional search.
m_A	$\uparrow$	A higher inherited minimum raises the breadth entering examination and, under the maintained minimum-dominance assumption, raises the value of examiner search.

Notes: The examiner chooses search to reduce expected overbreadth net of search cost. The sign for m_A uses the maintained minimum-dominance assumption: the increase in breadth entering examination dominates the induced increase in expected true breadth.

[4]The examiner payoff is intentionally modeled as depending only on expected overbreadth. This keeps the examiner's role focused on ex ante screening. Allowing examiner utility to depend on later challenge or invalidation outcomes, for example through reputational or career concerns, would add another incentive to search but would not change the core mechanism emphasized here.

6.5 Stagewise Comparative Statics of Applicant Search

I next characterize how applicant search responds to changes in the initial record, applicant incentives, and the incentives facing later players.

Applicant search is chosen before examination and challenge, so its comparative statics operate through the continuation value $Z(m_A, R_A; \theta_0)$. The relevant question is how each parameter changes the value of moving from the initial state (m_0, R_0) to the applicant-generated state (m_A, R_A).

The applicant's own search-cost shifter λ_A enters only through the cost function $c_A(s_A; \lambda_A)$. A higher λ_A raises the cost of higher search intensity and lowers the net return to search. Thus, the applicant's search objective exhibits decreasing differences in (s_A, λ_A).

The inherited record size R_0 affects the value of additional information. Holding m_0 fixed, a larger initial record leaves less uncertainty for applicant search to resolve and reduces the value of further search. Thus, the search objective exhibits decreasing differences in (s_A, R_0).

The inherited minimum m_0 affects search through the breadth entering the applicant's decision problem. A higher m_0 means the applicant begins from a less constraining initial record, increasing the value of learning whether broader protection is sustainable before examination. Under the maintained minimum-dominance assumption, the search objective exhibits increasing differences in (s_A, m_0).

Patent quality q affects applicant search through the continuation value of the patent. Higher quality raises the value of the eventual patent right and therefore raises the stakes of the information state carried into examination and challenge. Under the maintained quality-value dominance condition, the applicant's search objective exhibits increasing differences in (s_A, q).

The applicant's defense cost K_A lowers the continuation value of patents that are likely to face challenge. A higher K_A therefore reduces the value of applicant search as a way to manage downstream challenge risk, so the search objective exhibits decreasing differences in (s_A, K_A).

The remaining parameters operate through later-player behavior. A higher examiner search-cost shifter λ_E lowers examiner search and weakens expected truncation during prosecution, reducing the value of applicant search. A higher challenger search-cost shifter λ_C, or a higher challenge-cost environment, lowers downstream challenge pressure and therefore reduces the value of applicant search. Thus, the applicant's search objective exhibits decreasing differences in (s_A, λ_E), (s_A, λ_C), and the challenge-cost environment.

Combining these results, equilibrium applicant search is weakly increasing in q and m_0, and weakly decreasing in λ_A, K_A, R_0, λ_E, λ_C, and the challenge-cost environment. By Topkis's Theorem, there exists a monotone selection $s_A^*(\theta_0)$ with these properties [37].

Table 4 summarizes the stagewise comparative statics for applicant search.

Table 4: Stagewise comparative statics for applicant search

Parameter	Equilibrium search s_A^*	Main channel
λ_A	$\downarrow$	Higher applicant search costs lower the net return to search.
q	$\uparrow$	Higher patent quality raises the continuation value of the patent and increases the value of shaping the information state before examination.
K_A	$\downarrow$	Higher defense costs lower the continuation value of patents exposed to challenge risk.
λ_E	$\downarrow$	Higher examiner search costs reduce expected examination-stage scrutiny, lowering the value of applicant search.
λ_C	$\downarrow$	Higher challenger search costs reduce downstream challenge pressure, lowering the value of applicant search.
Higher challenge-cost environment	$\downarrow$	Higher challenge costs reduce expected filing pressure and therefore reduce the value of applicant search.
m_0	$\uparrow$	A higher inherited minimum raises the value of learning whether broader protection is sustainable.
R_0	$\downarrow$	A larger inherited record leaves less uncertainty for applicant search to resolve.

Notes: Applicant search maximizes the expected continuation value of the game net of search cost:

$$U_A(s_A;\theta_0) = \mathbb{E}\left[Z(m_A(s_A), R_A(s_A);\theta_0) \mid m_0, R_0\right] - c_A(s_A;\lambda_A).$$

The signs for q and m_0 use the maintained dominance conditions described in the text.

6.6 Global Monotone Comparative Statics

The stagewise comparative statics above characterize the direct responses within each player's decision problem. Global comparative statics are harder because a parameter change that moves applicant search also changes the record inherited by the examiner, which then induces an endogenous examiner response. To obtain sharp signs for the on-path prosecution record, I impose one additional restriction on how strongly examiner search offsets changes in applicant search.

The restriction is consistent with the institutional fact that patent examination is shaped by binding time allocations and workload pressure [12, 11, 18]. In the model, the same idea follows from diminishing returns to examiner search and increasing marginal search costs: examiner search may adjust in response to applicant search, but it does not fully undo it.

Assumption 4 (Partial Examiner Adjustment) *For any parameter change that alters equilibrium applicant search, the induced change in equilibrium examiner search is weakly in the opposite direction and strictly smaller in absolute value whenever applicant search changes. Formally, if a parameter change from τ to τ' changes equilibrium search from $s_A^*(\tau)$ to $s_A^*(\tau')$ and from $s_E^*(\tau)$ to $s_E^*(\tau')$, then*

$$\left(s_A^*(\tau') - s_A^*(\tau)\right)\left(s_E^*(\tau') - s_E^*(\tau)\right) \leq 0$$

and, whenever $s_A^*(\tau') \neq s_A^*(\tau)$,

$$\left|s_E^*(\tau') - s_E^*(\tau)\right| < \left|s_A^*(\tau') - s_A^*(\tau)\right|.$$

Let $\Delta x \equiv x(\tau') - x(\tau)$ denote the equilibrium change induced by a parameter shift. Assumption 4 implies that total pre-grant search moves in the same direction as applicant search:

$$\Delta s_A^* > 0 \quad \Longrightarrow \quad \Delta(s_A^* + s_E^*) > 0,$$

and

$$\Delta s_A^* < 0 \quad \Longrightarrow \quad \Delta(s_A^* + s_E^*) < 0.$$

Since the prosecution record is built from cumulative pre-grant search, the inherited record size R_E moves in the same direction as s_A^*. By the maintained information structure, higher cumulative search also shifts the effective minimum m_E in a more constraining direction, while lower cumulative search shifts it in a less constraining direction. The examiner therefore smooths the effect of applicant search on the prosecution record but does not overturn it.

The applicant's accept-or-abandon decision requires one further dominance restriction. A stronger prosecution record tends to improve the applicant's downstream position by reducing challenge risk and increasing expected merits survival, but it may also reduce granted breadth and therefore lower patent value. I maintain that, for equilibrium changes in the prosecution state, the downstream gains from a stronger record dominate any direct loss from induced narrowing. Under this continuation-value dominance restriction,

$$\Delta s_A^* > 0 \quad \Longrightarrow \quad \Delta \hat{K}_A > 0,$$

while

$$\Delta s_A^* < 0 \quad \Longrightarrow \quad \Delta \hat{K}_A < 0.$$

Thus, applicant acceptance inherits the direction of the prosecution record under the maintained restriction.

The challenger-side objects are less sharp. A parameter change can affect challenger search and filing through the challenger's own incentives while also changing applicant search and thereby altering the record the challenger later observes. In reduced form,

$$\tau \longrightarrow s_A^*(\tau) \longrightarrow (m_E(\tau), R_E(\tau)) \longrightarrow s_C^*(\tau),\, p_f(\tau).$$

The direct challenger-side effect and the indirect deterrence effect through the prosecution record may point in opposite directions. Without an additional dominance restriction on their relative magnitudes, the global effects on challenger search and filing probability remain ambiguous.

These observations can be summarized as follows.

Proposition 2 *Under Assumption 4 and the maintained dominance restriction for the applicant's continuation value, any parameter change that raises equilibrium applicant search raises cumulative pre-grant search, increases the prosecution record size, shifts the effective minimum in a more constraining direction, and weakly raises applicant acceptance. Any parameter change that lowers equilibrium applicant search has the opposite effects on the prosecution record and weakly lowers applicant acceptance. The global effects on challenger search and filing probability remain ambiguous in general.*

Proposition 2 gives the basic structure of the global comparative statics. The sharpest global predictions concern the prosecution-stage information state and the applicant-side equilibrium objects. Challenger-side outcomes remain more sensitive to offsetting equilibrium feedback because applicant behavior itself acts as an endogenous deterrence margin.

Table 5 collects the global comparative statics for the main equilibrium objects. These signs are the central empirical predictions of the model: they describe how changes in primitives should appear in applicant search, examiner scrutiny, granted breadth, acceptance, and post-grant challenge behavior.

Table 5: Global monotone comparative statics

Parameter	s_A^*	m_A	s_E^*	m_E	Acceptance	s_C^*	p_f
q	↑	↓	↓	↓	↑	?	?
λ_A	↓	↑	↑	↑	↓	?	?
λ_E	↓	↑	?	↑	↓	?	?
λ_C	↓	↑	↑	↑	↓	?	?
FOSD ↑ in G_C	↓	↑	↑	↑	↓	?	?
FOSD ↑ in G_A	↓	↑	↑	↑	↓	?	?
m_0	↑	↓	↓	↓	↑	?	?
R_0	↓	↑	↑	↑	↓	?	?

Notes: The signs summarize the global equilibrium effects in Proposition 2. A decrease in m_A or m_E means the relevant minimum shifts in a more constraining direction. Under Assumption 4, changes in applicant search move the prosecution-stage record in the same direction as cumulative pre-grant search. The sign on acceptance uses the maintained dominance restriction on the applicant's continuation value. Question marks indicate cases in which direct incentive effects and indirect equilibrium-record effects may offset one another. In the λ_E row, examiner search is ambiguous because higher examiner search costs lower s_E^* directly, while the induced reduction in applicant search can create an offsetting examiner response.

7 Empirical Content and Descriptive Evidence

This section connects the model to observable patent data. The model's equilibrium objects are not observed directly, but they correspond to standard measures of applicant disclosures, examiner-added information, claim breadth, prosecution outcomes, and post-grant challenges. These mappings generate testable predictions for how changes in search costs, challenge costs, and patent value should appear in the data.

7.1 From Model Objects to Empirical Proxies

The empirical content of the model turns on separating three margins: applicant-generated information, examiner-added information, and downstream challenge activity. Applicant and examiner search are not directly observed, and legally valid breadth b^* is latent. The goal is therefore not to find one perfect proxy for each model object, but to organize observable patent-data measures according to the sequence of the game.

Applicant search s_A is best interpreted as applicant-side information production. Its closest empirical counterparts are applicant-submitted prior-art references, IDS filings, applicant-submitted patent citations, and applicant-submitted non-patent literature. These variables measure information placed into the prosecution record rather than latent search effort itself, since U.S. law imposes a duty to disclose known material information but not a general duty to conduct a prior-art search. The empirical counterpart of R_A is therefore the size of the applicant-generated record observed by the examiner. Prior work emphasizes the importance of separating applicant-submitted and examiner-added references because they capture different parts of the prosecution process [1, 16, 7].

The applicant-stage minimum m_A corresponds to the breadth of the claims entering examination. Following the claim-breadth literature, this object can be proxied by first independent claim length, shortest independent claim length, or the number of independent claims at application or publication [28, 17]. The maintained interpretation is that, all else equal, shorter claims or more independent claiming routes tend to indicate broader breadth. Because claim-drafting conventions vary across technologies, these proxies should be used with technology fixed effects, art-unit controls, and robustness checks for fields where claim length may be less informative [28, 29].

Examiner search s_E maps into examiner-side information production and prosecution-stage scrutiny. Natural proxies include examiner-added citations, rejection art, office actions, Section 102 and Section 103 rejections, RCEs, appeals, interviews, and prosecution duration. Examiner-added citations are closest to the model's information-production margin, while office actions and rejections capture broader prosecution intensity. Prior work documents systematic variation across examiners and institutional environments, supporting the interpretation of examination as an endogenous screening margin [5, 22, 12, 11, 10].

Granted breadth $b_E = m_E$ corresponds to claim breadth at grant. In the data, m_E can be proxied using grant-stage claim-breadth measures such as first independent claim length, shortest independent claim length, or independent claim count. The difference between application-stage and grant-stage breadth provides a measure of prosecution-stage narrowing [28, 4, 29]. This narrowing reflects both examiner objections and applicant amendment choices, so it should be interpreted as an equilibrium prosecution outcome rather than a pure measure of examiner effort.

The applicant accept-or-abandon decision maps into whether the application continues through prosecution and ultimately issues as a patent. Observable outcomes include grant, abandonment, continuation filings, continuation-in-part filings, divisional applications, RCEs, and appeals. These outcomes are useful because the model predicts that search and challenge risk affect not only granted breadth but also the applicant's willingness to continue pursuing protection.

The challenger stage is observed through post-grant activity. Filing Y_C and filing probability p_f map into district court litigation, IPR petitions, PGR petitions, institution decisions, final written decisions, and claim cancellation or amendment. Challenger search s_C is harder to observe, but asserted prior art in PTAB petitions or litigation provides a partial measure of challenger-side information production. Post-grant proceedings are especially useful because they reveal whether challengers rely on art already in the prosecution record or instead identify new prior art not previously considered [27, 3, 44].

Table 6 summarizes the main empirical counterparts. The mapping is not meant to imply that any single variable perfectly measures the corresponding theoretical object. Rather, the table shows how the model organizes standard patent-data objects into a coherent sequence: applicant-generated information, examiner-added information, prosecution-stage narrowing, applicant continuation, and downstream challenge.

This mapping clarifies why post-grant review reforms should affect more than observed challenge filings. In the model, a change in the downstream challenge environment can alter applicant search, examiner scrutiny, claim breadth, and grant or abandonment decisions before any challenge is filed. The empirical analysis should therefore examine both downstream outcomes, such as PTAB filings and litigation, and upstream margins, such as applicant-submitted prior art, examiner-added prior art, claim narrowing, and grant outcomes.

7.2 Descriptive Facts Motivating the Empirical Margins

Table 6 maps the model's main objects into observable patent-data counterparts. The next question is whether these margins are empirically important. This subsection highlights three descriptive facts from the existing literature. These facts are not direct tests of the model; rather, they motivate the empirical relevance of treating claim breadth, applicant-generated information, examiner-generated information, and post-grant challenge as distinct but connected margins.

First, downstream challenges are selective. If all granted patents faced the same post-grant threat, the challenger stage would have little empirical content. In practice, challenges are concentrated on patents with high stakes and high expected returns to invalidation. Prior work on IPR and related proceedings shows challenged patents differ systematically from unchallenged patents along measures of breadth, value, and legal vulnerability [27, 3]. Martinelli and Mazzei provide closely related evidence using claim-breadth measures, showing that patents subject to post-grant opposition tend

Table 6: Model objects and empirical counterparts

Model object	Interpretation	Empirical counterparts
s_A	Applicant search / applicant information production	Applicant-submitted citations; IDS filings; applicant patent citations; applicant non-patent literature citations; references submitted before or during prosecution.
R_A	Size of the applicant-generated record	Number of applicant-submitted references; number of IDS filings; applicant-submitted patent and non-patent references.
m_A	Breadth carried into examination	First independent claim length at application or publication; shortest independent claim length; independent claim count; application-stage claim-breadth measures.
s_E	Examiner search and scrutiny	Examiner-added citations; examiner-used rejection art; office actions; Section 102 and Section 103 rejections; interviews; appeals; RCEs; prosecution duration.
$R_E - R_A$	Examiner-added record	Examiner-added references; examiner citations used in rejection; additional prior art introduced during prosecution.
$m_E = b_E$	Granted breadth	First independent claim length at grant; shortest independent claim length at grant; independent claim count at grant; grant-stage claim-breadth measures.
$m_A - m_E$ or Δm	Prosecution-stage narrowing	Change in claim length from application or publication to grant; amendments to independent claims; reduction in independent claim count; difference between application and grant breadth.
Y_A	Applicant acceptance / continuation	Grant; abandonment; continuation filings; continuation-in-part filings; divisional filings; RCEs; appeals.
s_C	Challenger search	Prior art asserted in IPR, PGR, or litigation; newly introduced references in post-grant proceedings; overlap between asserted art and prosecution-history art.
Y_C and p_f	Challenge filing and challenge probability	IPR petitions; PGR petitions; district court litigation; institution decisions; final written decisions; claim cancellation or amendment.
q	Patent value / stakes of breadth	Forward citations; family size; PCT filing; renewal or maintenance behavior; assignment or transaction measures.
λ_A	Applicant search-cost shifter	Applicant sophistication; attorney characteristics; prior-art complexity; technology field; firm size; experience with patent prosecution.
λ_E	Examiner search-cost shifter	Examiner workload; examiner seniority; art-unit backlog; time allocations; technology-specific prior-art complexity; institutional examination constraints.
λ_C	Challenger search-cost shifter	Litigation or PTAB search costs; availability of prior-art search tools; technology-specific prior-art complexity; party sophistication; challenger resources.
K_C and K_A	Challenge and defense costs	AIA/PTAB availability; forum; technology-level litigation cost; party type; entity size; NPE involvement; expected cost of challenge.

to have broader claims than otherwise similar non-opposed patents [29]. This fact corresponds to the filing condition in the model: broader and more valuable patents raise the return to challenge and may increase expected success by increasing expected overbreadth.

Second, post-grant challenge institutions can affect behavior before any challenge is filed. This is the upstream discipline mechanism in the model. A change in the cost or credibility of challenge changes the expected payoff from seeking broad claims and can therefore alter applicant and examiner behavior during prosecution. Martinelli and Mazzei study the introduction of post-grant review through the AIA using U.S.-Canada patent twins and find that applicants respond by filing narrower claims at the USPTO after the reform. They also find greater prosecution-stage narrowing when applicants do not preemptively narrow claims themselves, consistent with examiners partly offsetting applicant inaction [29]. In the language of the model, a change in the downstream challenge environment shifts not only Y_C or p_f, but also upstream objects such as m_A, m_E, and the prosecution record inherited by later players.

Third, applicant-submitted information and examiner-generated information are empirically distinct. The model treats applicant search and examiner search as separate margins because the applicant-generated record is observed by the examiner and then augmented during examination. Patent-citation evidence supports this separation. Applicant-submitted and examiner-added references differ in timing, source, and empirical interpretation [1]. Prior work also shows that applicant and examiner citations have different relationships with patent value and prosecution outcomes [16, 7]. This distinction matters because total prior art in the record is not enough to identify the mechanism: a richer prosecution record may reflect applicant search, examiner search, or both.

Together, these facts suggest that the empirical content of the model lies in the joint behavior of multiple margins. Claim breadth is related to downstream challenge, post-grant review can feed back into upstream drafting and examination, and the prosecution record contains distinct applicant and examiner components. A full empirical analysis should therefore study applicant-side information, examiner-side information, claim narrowing, grant or abandonment, and post-grant challenge jointly. The purpose of the present paper is to provide a framework explaining why these margins should be connected.

7.3 Testable Predictions

The preceding subsections map the model into observable patent-data objects and show that these margins are empirically meaningful. I now use the comparative statics to organize a set of testable predictions. The goal is not to estimate these relationships in the present paper, but to clarify what an empirical implementation of the theory would look for.

A natural empirical setting is the introduction of post-grant review through the America Invents Act. In the language of the model, the AIA can be interpreted primarily as a shift toward a lower challenge-cost environment, since it created new administrative procedures for challenging issued patents through the PTAB. This interpretation has two caveats. First, the AIA was not a single-parameter change; it also included other institutional changes, most notably the transition from first-to-invent to first-to-file. Second, the relevant timing may differ across players. Applicants may respond when the legal change becomes anticipated, while examiners may respond more strongly once post-grant review becomes effective and their examination outcomes can be scrutinized through the new procedures [29]. The AIA is therefore best viewed as a useful setting for studying changes in challenge risk, not as a perfect experiment that moves only one primitive of the model.

The predictions should be interpreted as weak directional predictions. When the model predicts an increase or decrease in an equilibrium object, the empirical implication is that the object should move in the predicted direction or display no meaningful response. Some margins may appear close to zero because the relevant constraint does not bind, the treatment is weak for patents unlikely to face challenge, or offsetting equilibrium channels make the net response small. The empirical content of the model is therefore not that every margin responds strongly in every sample, but that systematic responses should align with the directional logic of the comparative statics.

The first set of predictions concerns applicant behavior. When downstream challenge becomes more severe, applicants have stronger incentives to learn about the limits of valid protection before examination. In the model, this appears as higher applicant search s_A and narrower breadth entering examination, m_A. Empirically, this predicts weak increases in applicant-side information measures, such as applicant-submitted citations, IDS filings, applicant patent citations, and applicant non-patent literature citations. It also predicts weak narrowing of application-stage claims, measured by longer first independent claims, longer shortest independent claims, or fewer independent claiming routes.

The second set of predictions concerns examiner scrutiny and the prosecution record. Because applicant and examiner search interact, a stronger applicant response can reduce the marginal value of examiner search. Under the partial-adjustment condition used for the global comparative statics, lower challenge costs increase applicant search and strengthen the applicant-generated record, while examiner search may fall in response. Nevertheless, total pre-grant information increases and granted breadth falls. Empirically, this means examiner-added citations or office actions should be interpreted jointly with the combined prosecution record and the amount of narrowing between application and grant. A weak decline in examiner-side search measures need not imply weaker screening if it occurs alongside a stronger applicant-generated record and narrower granted claims.

The third set of predictions concerns grant outcomes and post-grant challenge. The model predicts that a stronger prosecution record can increase the applicant's willingness to accept the resulting patent because the patent is less vulnerable conditional on challenge. Grant, abandonment, continuation, and RCE outcomes are therefore relevant applicant-side margins. The prediction for observed post-grant challenges is less sharp. A lower challenge-cost environment directly makes filing more attractive to challengers, but it also induces upstream responses that may make granted patents less vulnerable. Thus, PTAB availability need not mechanically increase observed challenge rates. The observed effect on IPR petitions, PGR petitions, litigation, institution, or final written decisions depends on the relative strength of the direct challenge-cost effect and the indirect upstream-discipline effect.

Taken together, these predictions imply that an empirical implementation should examine the full sequence of patenting behavior rather than only the final challenge outcome. A post-grant review reform may affect applicant-side information production, examiner-side scrutiny, claim breadth, grant or abandonment, and post-grant challenge at the same time. The relevant empirical question is whether the collection of responses across these margins is consistent with downstream challenge risk feeding back into the production of patent breadth before grant.

7.4 Empirical Implementation in Future Work

The empirical implementation of these predictions is left for future work. The natural next step is to combine patent text, prosecution-history data, citation-source information, and post-grant challenge outcomes in order to follow the full sequence emphasized by the model. The analysis should examine not only whether a patent is eventually challenged, but also whether changes in the downstream challenge environment affect applicant-generated information, examiner-added information, claim breadth at application, prosecution-stage narrowing, grant or abandonment, and post-grant challenge behavior.

One empirical setting is the introduction of post-grant review through the AIA. A simple design would compare patents more and less exposed to post-grant challenge risk before and after the reform. Exposure could be measured using pre-reform technology-level litigation or challenge rates, patent value, claim breadth, patent-thicket density, or other predictors of downstream challenge. The basic specification could relate an outcome Y_{it} to the interaction between the post-AIA period and predicted challenge exposure:

$$Y_{it} = \beta \left(\text{PostAIA}_t \times \text{Threat}_i\right) + X_{it}'\Gamma + \alpha_{j(i)} + \lambda_t + \varepsilon_{it},$$

where Y_{it} may measure applicant-submitted prior art, examiner-added prior art, application claim breadth, prosecution-stage narrowing, grant or abandonment, or post-grant challenge. Technology fixed effects $\alpha_{j(i)}$ and time fixed effects λ_t absorb persistent differences across technological areas and aggregate changes in patenting behavior.

A related design could follow the patent-twins approach used in prior work and compare U.S. applications to foreign counterparts for the same underlying invention [29]. Family fixed effects would absorb invention-level heterogeneity, with identifying variation coming from whether the U.S. member of the family is exposed to the post-AIA challenge environment. A richer version of this design could interact U.S. exposure after the reform with predicted challenge threat:

$$Y_{ift} = \beta \left(\text{US}_i \times \text{PostAIA}_t \times \text{Threat}_f\right) + \alpha_f + \lambda_t + \varepsilon_{ift},$$

where α_f denotes patent-family fixed effects. This design is especially useful for studying claim breadth and prosecution-stage narrowing because it compares the same invention across legal environments.

The purpose of the present paper is not to carry out this empirical exercise, but to discipline what such an exercise should measure. The model implies that post-grant review should not be evaluated only by counting post-grant challenges. A reform that lowers challenge costs may also change the information produced by applicants, the information added by examiners, the breadth of claims that enter and leave prosecution, and the set of patents applicants choose to continue pursuing. A full empirical analysis should therefore treat the patenting process as a sequence of linked margins rather than as a single post-grant outcome.

8 Conclusion

Patent breadth is often discussed as if it were determined either by the applicant's initial claim-drafting decision or by the patent office's grant decision. Instead, I argue that patent breadth is the equilibrium outcome of a sequential search process involving applicants, examiners, and potential challengers. Applicants decide how much information to acquire before examination, examiners decide how much additional information to acquire during prosecution, and challengers decide whether further search and filing are worthwhile after grant. The breadth that ultimately matters is therefore shaped by strategic information acquisition across the entire patenting sequence.

The model formalizes this idea in a three-player dynamic game. The central state variable is the prosecution record, summarized by (m, R), where m is the most constraining draw and R is the size of the accumulated record. This structure captures two roles of search. Search can uncover a more constraining piece of prior art, lowering the breadth that can be supported, and it can enlarge the record, changing how later players form beliefs about the legally valid breadth of the patent. These two channels link applicant search, examiner search, and challenger search through a common information structure.

The main implication is that the stages of the patenting process cannot be understood independently. Applicant search shapes the record inherited by the examiner and later by the challenger. Examiner search affects both the breadth that survives prosecution and the public record on which downstream challenge decisions are based. Challenger search and filing occur after grant, but the expected threat of challenge feeds back into applicant and examiner incentives before grant. As a result, post-grant institutions can alter upstream behavior even before any challenge is filed.

The comparative statics highlight this equilibrium feedback. The sharpest predictions concern the applicant-generated record, the prosecution-stage information state, and the breadth that survives examination. Parameters that raise applicant search tend to strengthen the prosecution record and shift granted breadth in a more constrained direction, provided examiner adjustment is only partial. By contrast, the effect on observed post-grant challenge is generally ambiguous. Lower challenge costs directly make filing more attractive to challengers, but they can also induce upstream responses that make issued patents less vulnerable. Observed filing therefore depends on the balance between the direct challenge-cost effect and the indirect deterrence effect operating through applicant and examiner behavior.

This logic has implications for patent policy. Post-grant review should not be evaluated only by counting challenges filed or claims invalidated after grant. Those outcomes are important, but they are the final part of a longer strategic sequence. A reform that changes the cost or credibility of challenge may also change applicant search, the information placed into the prosecution record, examiner scrutiny, claim narrowing, and the set of applications applicants choose to continue pursuing. Similarly, policies that affect examiner capacity or applicant search incentives can have downstream effects on challenge behavior because they change the record later challengers inherit. The relevant policy environment is therefore the joint system of pre-grant and post-grant incentives.

The model also clarifies why patent breadth is a useful object for studying patent quality. A binary distinction between valid and invalid patents misses much of the economic action. Many patents matter because of how far their claims extend, how much uncertainty they create for rivals, and how costly they are to challenge. By focusing on breadth, the framework connects prior-art search, examination, and post-grant challenge through a common object.

Future empirical work can take these predictions to data by linking patent text, prosecution histories, citation-source information, and post-grant challenge outcomes. The relevant empirical object is not only whether a patent is eventually challenged, but the full sequence of margins through which challenge risk can affect patenting: applicant-generated information, examiner-added information, claim breadth at application and grant, grant or abandonment, and post-grant challenge. A companion empirical analysis can use changes in post-grant challenge institutions, such as the introduction of PTAB review, to study whether downstream challenge risk feeds back into the production of patent breadth before grant.

References

[1] Juan Alcácer, Michelle Gittelman, and Bhaven Sampat. "Applicant and examiner citations in U.S. patents: An overview and analysis". In: *Research Policy* 38.2 (2009), pp. 415–427. ISSN:

0048-7333. DOI: 10.1016/j.respol.2008.12.001. URL: https://www.sciencedirect.com/science/article/pii/S0048733308002825.

[2] Vidya Atal and Talia Bar. "Prior art: To search or not to search". In: *International Journal of Industrial Organization* 28.5 (2010), pp. 507–521. ISSN: 0167-7187. DOI: 10.1016/j.ijindorg.2009.12.002. URL: https://www.sciencedirect.com/science/article/pii/S0167718710000020.

[3] Talia Bar and Brendan Costello. "Predicting Administrative Patent Challenges". In: *Minn. J. Law Sci. Technol.* 21.1 (2020), p. 207.

[4] Rudi Bekkers, Arianna Martinelli, and Federico Tamagni. "The impact of including standards-related documentation in patent prior art: Evidence from an EPO policy change". In: *Research Policy* 49.7 (None 2020), None. DOI: 10.1016/j.respol.2020.104007. URL: https://ideas.repec.org/a/eee/respol/v49y2020i7s004873332030086x.html.

[5] Iain M. Cockburn, Samuel Kortum, and Scott Stern. *Are All Patent Examiners Equal? The Impact of Examiner Characteristics*. Working Paper 8980. National Bureau of Economic Research, June 2002. DOI: 10.3386/w8980. URL: http://www.nber.org/papers/w8980.

[6] Christopher A. Cotropia. "Modernizing Patent Law's Inequitable Conduct Doctrine". In: *Berkeley Technology Law Journal* 24.2 (Mar. 2009), p. 723. DOI: 10.15779/Z38T99F. URL: https://lawcat.berkeley.edu/record/1122301.

[7] Christopher A. Cotropia, Mark A. Lemley, and Bhaven Sampat. "Do applicant patent citations matter?" In: *Research Policy* 42.4 (2013), pp. 844–854. ISSN: 0048-7333. DOI: 10.1016/j.respol.2013.01.003. URL: https://www.sciencedirect.com/science/article/pii/S0048733313000085.

[8] *Exergen Corp. v. Wal-Mart Stores, Inc.* 575 F.3d 1312 (Fed. Cir. 2009). United States Court of Appeals for the Federal Circuit. Aug. 2009. URL: https://law.justia.com/cases/federal/appellate-courts/cafc/06-1491/06-1491-2011-03-27.html.

[9] Joseph Farrell and Carl Shapiro. "How Strong Are Weak Patents?" In: *The American Economic Review* 98.4 (2008), pp. 1347–1369. ISSN: 00028282. URL: http://www.jstor.org/stable/29730125 (visited on 04/19/2026).

[10] Josh Feng and Xavier Jaravel. "Crafting Intellectual Property Rights: Implications for Patent Assertion Entities, Litigation, and Innovation". In: *American Economic Journal: Applied Economics* 12.1 (Jan. 2020), pp. 140–181. DOI: 10.1257/app.20180361. URL: https://www.aeaweb.org/articles?id=10.1257/app.20180361.

[11] Michael D. Frakes and Melissa F. Wasserman. "Investing in Ex Ante Regulation: Evidence from Pharmaceutical Patent Examination". In: *American Economic Journal: Economic Policy* 15.3 (Aug. 2023), pp. 151–183. DOI: 10.1257/pol.20200703. URL: https://www.aeaweb.org/articles?id=10.1257/pol.20200703.

[12] Michael D. Frakes and Melissa F. Wasserman. "Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents?: Evidence from Micro-Level Application Data". In: *The Review of Economics and Statistics* 99.3 (2017), pp. 550–563. ISSN: 00346535, 15309142. URL: https://www.jstor.org/stable/26616139 (visited on 04/19/2026).

[13] Jerry R. Green and Suzanne Scotchmer. "On the Division of Profit in Sequential Innovation". In: *RAND Journal of Economics* 26.1 (1995), pp. 20–33. URL: https://ideas.repec.org/a/rje/randje/v26y1995ispringp20-33.html.

[14] Bronwyn H. Hall, Adam Jaffe, and Manuel Trajtenberg. "Market Value and Patent Citations". In: *The RAND Journal of Economics* 36.1 (2005), pp. 16–38. ISSN: 07416261. URL: http://www.jstor.org/stable/1593752 (visited on 04/19/2026).

[15] Dietmar Harhoff, Frederic M. Scherer, and Katrin Vopel. "Citations, family size, opposition and the value of patent rights". In: *Research Policy* 32.8 (2003), pp. 1343–1363. ISSN: 0048-7333. DOI: 10.1016/S0048-7333(02)00124-5. URL: https://www.sciencedirect.com/science/article/pii/S0048733302001245.

[16] Deepak Hegde and Bhaven Sampat. "Examiner citations, applicant citations, and the private value of patents". In: *Economics Letters* 105.3 (2009), pp. 287–289. ISSN: 0165-1765. DOI: 10.1016/j.econlet.2009.08.019. URL: https://www.sciencedirect.com/science/article/pii/S0165176509002845.

[17] Neil Thompson Jeffrey M. Kuhn. "How to Measure and Draw Causal Inferences with Patent Scope". In: *SSRN* (2019). URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2977273.

[18] Yee Kyoung Kim and Jun Byoung Oh. “Examination workloads, grant decision bias and examination quality of patent office”. In: *Research Policy* 46.5 (2017), pp. 1005–1019. ISSN: 0048-7333. DOI: 10.1016/j.respol.2017.03.007. URL: https://www.sciencedirect.com/science/article/pii/S0048733317300586.

[19] Paul Klemperer. “How Broad Should the Scope of Patent Protection Be?” In: *RAND Journal of Economics* 21.1 (1990), pp. 113–130. URL: https://ideas.repec.org/a/rje/randje/v21y1990ispringp113-130.html.

[20] Corinne Langinier and Philippe Marcoul. “Monetary and implicit incentives of patent examiners”. In: *Journal of Economics and Business* 110.C (2020). DOI: 10.1016/j.jeconbus.2020.105906. URL: https://ideas.repec.org/a/eee/jebusi/v110y2020ics0148619519303078.html.

[21] Corinne Langinier and Philippe Marcoul. “The Search of Prior Art and the Revelation of Information by Patent Applicants”. In: *Review of Industrial Organization* 49.3 (2016), pp. 399–427. ISSN: 0889938X, 15737160. URL: http://www.jstor.org/stable/44735431 (visited on 03/24/2026).

[22] Mark A. Lemley and Bhaven Sampat. “Examiner Characteristics and Patent Office Outcomes”. In: *The Review of Economics and Statistics* 94.3 (2012), pp. 817–827. ISSN: 00346535, 15309142. URL: http://www.jstor.org/stable/23261480 (visited on 04/19/2026).

[23] Mark A. Lemley and Carl Shapiro. “Probabilistic Patents”. In: *Journal of Economic Perspectives* 19.2 (June 2005), pp. 75–98. DOI: 10.1257/0895330054048650. URL: https://www.aeaweb.org/articles?id=10.1257/0895330054048650.

[24] Joshua Lerner. “The Importance of Patent Scope: An Empirical Analysis”. In: *The RAND Journal of Economics* 25.2 (1994), pp. 319–333. ISSN: 07416261. URL: http://www.jstor.org/stable/2555833 (visited on 04/19/2026).

[25] Chol-Won Li. “On the Policy Implications of Endogenous Technological Progress”. In: *The Economic Journal* 111.471 (Apr. 2001), pp. C164–C179. ISSN: 0013-0133. DOI: 10.1111/1468-0297.00626. eprint: https://academic.oup.com/ej/article-pdf/111/471/C164/26553530/ejc164.pdf. URL: https://doi.org/10.1111/1468-0297.00626.

[26] Brian J. Love and Shawn Ambwani. "Inter Partes Review: An Early Look at the Numbers". In: *The University of Chicago Law Review Dialogue* 81 (Oct. 2014). URL: https://chicagounbound.uchicago.edu/uclrev_online/vol81/iss1/5/.

[27] Brian J. Love, Shawn P. Miller, and Shawn Ambwani. "Determinants of Patent Quality: Evidence from Inter Partes Review Proceedings". In: *University of Colorado Law Review* 90 (2019), p. 67. URL: https://scholar.law.colorado.edu/lawreview/vol90/iss1/3/.

[28] Alan C. Marco, Joshua D. Sarnoff, and Charles A. W. deGrazia. "Patent claims and patent scope". In: *Research Policy* 48.9 (2019), pp. 1–1. DOI: 10.1016/j.respol.2019.04.014. URL: https://ideas.repec.org/a/eee/respol/v48y2019i926.html.

[29] Arianna Martinelli and Julia Mazzei. "Death squad or quality improvement? The impact of introducing post-grant review on U.S. patent legal quality". In: *Research Policy* 54.5 (2025), p. 105205. ISSN: 0048-7333. DOI: 10.1016/j.respol.2025.105205. URL: https://www.sciencedirect.com/science/article/pii/S0048733325000344.

[30] Paul Milgrom and Chris Shannon. "Monotone Comparative Statics". In: *Econometrica* 62.1 (1994), pp. 157–180. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/2951479 (visited on 06/03/2026).

[31] Ted O'Donoghue. "A Patentability Requirement for Sequential Innovation". In: *RAND Journal of Economics* 29.4 (Winter 1998), pp. 654–679. DOI: None. URL: https://ideas.repec.org/a/rje/randje/v29y1998iwinterp654-679.html.

[32] Ted O'Donoghue and Josef Zweimüller. "Patents in a Model of Endogenous Growth". In: *Journal of Economic Growth* 9.1 (2004), pp. 81–123. ISSN: 13814338, 15737020. URL: http://www.jstor.org/stable/40215894 (visited on 06/02/2026).

[33] Mark Schankerman and Florian Schuett. "Patent Screening, Innovation, and Welfare". In: *The Review of Economic Studies* 89.4 (2021), pp. 2101–2148. ISSN: 0034-6527. DOI: 10.1093/restud/rdab073. URL: https://doi.org/10.1093/restud/rdab073.

[34] Florian Schuett. "Patent quality and incentives at the patent office". In: *The RAND Journal of Economics* 44.2 (2013), pp. 313–336. DOI: 10.1111/1756-2171.12021. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12021.

[35] Suzanne Scotchmer. “Standing on the Shoulders of Giants: Cumulative Research and the Patent Law”. In: *Journal of Economic Perspectives* 5.1 (Mar. 1991), pp. 29–41. DOI: 10.1257/jep.5.1.29. URL: https://www.aeaweb.org/articles?id=10.1257/jep.5.1.29.

[36] *Therasense, Inc. v. Becton, Dickinson & Co.* 649 F.3d 1276 (Fed. Cir. 2011) (en banc). United States Court of Appeals for the Federal Circuit. 2011. URL: https://law.justia.com/cases/federal/appellate-courts/cafc/08-1511/08-1511-2011-03-27.html.

[37] Donald M. Topkis. “Minimizing a Submodular Function on a Lattice”. In: *Operations Research* 26.2 (1978), pp. 305–321. ISSN: 0030364X, 15265463. URL: http://www.jstor.org/stable/169636 (visited on 03/23/2026).

[38] United States Patent and Trademark Office. *America Invents Act: Inter Partes Disputes*. Accessed March 23, 2026. 2026. URL: https://www.uspto.gov/patents/laws/america-invents-act-aia/inter-partes-disputes.

[39] United States Patent and Trademark Office. *Information Disclosure Statement (IDS) Resources*. Accessed March 23, 2026. 2026. URL: https://www.uspto.gov/patents/apply/filing-online/information-disclosure-statement-ids-resources.

[40] United States Patent and Trademark Office. *Manual of Patent Examining Procedure, Section 2001: Duty of Disclosure, Candor, and Good Faith*. Accessed March 23, 2026. 2026. URL: https://www.uspto.gov/web/offices/pac/mpep/s2001.html.

[41] United States Patent and Trademark Office. *Manual of Patent Examining Procedure, Section 704: Nature of Examiner's Action*. Accessed March 23, 2026. 2026. URL: https://www.uspto.gov/web/offices/pac/mpep/s704.html.

[42] United States Patent and Trademark Office. *Manual of Patent Examining Procedure, Section 713: Interviews*. Accessed March 23, 2026. 2026. URL: https://www.uspto.gov/web/offices/pac/mpep/s713.html.

[43] United States Patent and Trademark Office. *Patent Process Overview*. Accessed March 23, 2026. 2026. URL: https://www.uspto.gov/patents/basics/patent-process-overview.

[44] Stephen Yelderman. “Prior Art in the District Court”. In: *Notre Dame L. Rev.* 95 (2019), p. 837. URL: https://scholarship.law.nd.edu/law_faculty_scholarship/1388.