Standard Setting Committees *

Tim Simcoe
J.L. Rotman School of Management, University of Toronto

December 2008

*Financial support for this research was provided by the UC Berkeley center for Information Technology Research in the Interest of Society, and the Intel Robert Noyce Memorial Scholarship. Special thanks to David Mowery, Paul Gertler, Hal Varian, and Bronwyn Hall and Stephanie Gates. I received helpful comments from three anonymous referees, Pierre Azoulay, Meghan Busse, Scott Bradner, Ken Corts, Rui deFigueiredo, Alberto Galasso, Ig Horstmann, Catherine McArthy, Dov Rothman, Marc Rysman, Dan Snow, Jason Snyder, David Waguespack, Catherine Wolfram, and participants in the Business and Public Policy Student Seminar at UC Berkeley. Address for correspondence: Joseph L. Rotman School of Management, 105 St. George Street, Toronto, ON M5S 3E6,Canada. E-mail: timothy.simcoe@rotman.utoronto.ca

Standard Setting Committees

Abstract

Voluntary Standard Setting Organizations (SSOs) use a consensus process to create new compatibility standards. Practitioners have suggested that SSOs are increasingly "politicized" and perhaps incapable of producing timely standards. This paper develops a simple model of standard setting committees and tests its predictions using data from the Internet Engineering Task Force, an SSO that produces many of the standards used to run the Internet. The results show that an observed slowdown in standards production between 1993 and 2003 can be linked to distributional conflicts created by the rapid commercialization of the Internet.

1 Introduction

The economic literature on compatibility often stresses standards wars, such as the recent video-format battle between Blu-ray and HD-DVD. However, many important standards emerge from voluntary Standard Setting Organizations (SSOs), which seek a broad consensus among interested parties before endorsing a particular technology and promoting its adoption. Shapiro & Varian (1998, 240) describe formal standard setting as “a wild mix of politics and economics.” This paper asks when that process is likely to work well.

I begin with a simple model that shows how consensus decision-making can produce lengthy delays. In this model, delay can be efficient, since it leads to better outcomes. However, rent-seeking produces excessive delay when participants favor specific technologies because of development lead times, proprietary complements, or intellectual property rights. As the private benefits of adopting a preferred alternative increase, coordination delays grow longer and less efficient.

To test the model, I collect detailed committee and proposal-level data from the Internet Engineering Task Force (IETF), an influential SSO that produces many of the standards used to run the Internet. These data cover a period between 1993 and 2003 when rapid Internet commercialization led to significant changes in the size and demographics of the IETF. I use these changes to construct a committee-level measure of commercial relevance — the “suit-to-beard” ratio — that proxies for participants’ private interest in specific outcomes.

To estimate the impact of rent-seeking on time-to-consensus, I develop a difference-in-differences estimator that exploits a unique feature of the IETF standards process. As described below, the IETF uses “nonstandards” to publish new ideas without formally endorsing them. Since standards and nonstandards go through an identical publication process, I treat these documents as a no-conflict control sample and use them to estimate the relevant counterfactual, i.e. how long it would take a committee to reach consensus in the absence of competing interests.

The empirical results show a statistically and economically significant correlation between distributional conflict and slower standards production. Specifically, a one percentage-point increase in private-sector participation (i.e. the “suit-to-beard” ratio) adds between 5 and 8 days to the standards development process. Since private-sector participation in IETF committees grew by 30 percentage points during the 1990s, these estimates suggest that Internet commercialization led to an additional seven months of deliberation for a typical standard. This effect has grown over time, and is larger for committees working on standards that are more likely to interface with proprietary technology. I also provide evidence that longer delays are associated with greater *ex post* impact, as measured by citations.

Overall, the empirical results suggest that rapid Internet commercialization led to an increase in strategic maneuvering within the IETF and a slowdown in committee decision-making. More broadly, they point to rent-seeking as a source of coordination costs, which many economists (e.g. Becker & Murphy, 1992) argue are a key factor limiting the potential gains from specialization and the division of labor.

1.1 Related Literature

While firms often design products to work together, inter-operability can be challenging for large systems with many independently supplied components. Competition offers one route to coordination. However, the literature on standards wars suggests that it will be especially intense and uncertain when network effects are strong (see, for example, David & Greenstein (1990) or Besen & Farrell (1994)). Another possibility is to allow large "platform leaders" to orchestrate major technical transitions. But dominant platforms need not have dominant firms; as emphasized by Bresnahan & Greenstein (1999) in their historical account of the computer industry. A third path to inter-operability is for firms with a stake in the platform to create an institution for collective self-governance. In the information and communications technology sector, these groups are called Standard Setting Organizations.[1]

SSOs seek a broad consensus before endorsing a particular technology. Since they are voluntary organizations, and typically lack enforcement power, we might expect these recommendations to have little impact. But network effects may cause cues for coordination to become self-enforcing once a consensus has been reached. Rysman & Simcoe (2008) provide indirect evidence of this "endorsement effect" by documenting an increase in citations to patents after they are disclosed to an SSO.[2] This paper emphasizes a related idea: if endorsement generates substantial rents, firms may fight for a long time before reaching consensus.

In fact, practitioners often point to delay as a source of significant opportunity costs and a major problem with the formal standards process (see National Research Council (1990) or Cargill (2001)). Moreover, a rise in standards-related patent disputes; the proliferation of industry-sponsored consortia; and several antitrust actions have led observers, such as Lemley (2007) and Updegrove (2007), to suggest that SSOs are becoming more "politicized" and perhaps less capable of producing timely standards.[3]

[1] The "ICT Consortia List" compiled by the European Committee for Standardization (www.cenorm.be) identifies 298 different Standards Setting Organizations and www.consortiuminfo.org lists more than 400 (accessed December 15, 2006).

[2] The law and economics literature generally *assumes* that an SSO endorsement leads to large sunk investments and considers the problem of *ex post* hold-up by patent-owners. SSOs try to mitigate this problem by requiring participants to disclose patents and license essential technology on "reasonable and non-discriminatory" (RAND) terms. For more on this issue, see Lemley (2002) or Farrell *et al.* (2007).

[3] Recent antitrust cases include *Dell Computer Corporation*, 121 F.T.C. 616; *Rambus Incorporated*, F.T.C Docket 9302; and *Negotiated Data Solutions LLC*, F.T.C. File No. 0510094.

Farrell & Saloner (1988) use the war of attrition to model coordination delays in the formal standards process.[4] Their theory predicts that SSOs are slower than markets, but more likely to achieve coordination. Farrell & Simcoe (2008) add private information to the war of attrition and ask when the benefits of screening for better technology are worth the costs of delay. Lerner & Tirole (2006) develop an alternative theory of SSOs that stresses their role as a certification agent rather than a forum for reaching consensus. In their model, technology vendors choose the friendliest SSO whose certification will persuade end-users to adopt the standard. Chiao *et al.* (2007) analyze the bylaws and intellectual property policies of fifty-nine SSOs and find support for this hypothesis.

While much of the early empirical research on SSOs was descriptive (e.g. Besen & Johnson (1988); Besen & Saloner (1989); Farrell & Shapiro (1992); Foray (1995); Lehr (1995)), recent studies have quantified various aspects of formal standardization. For example, Rosenkopf *et al.* (2001) study alliance formation among firms that participate in the same standards committee. Bekkers *et al.* (2002) show how SSO members' patent portfolios and alliance networks evolve over time. And Waguespack & Fleming (2008) show that IETF participation is correlated with liquidity events in a sample of start-ups. This paper is perhaps closest to Weiss & Sirbu (1990) or Leiponen (2008), which use administrative data from SSOs to study the the determinants of successful proposals.

The balance of the paper is organized into five parts. Section 2 contains a simple model of standard setting committees. Section 3 describes the IETF. Section 4 presents the empirical strategy, and Section 5 discusses results. Section 6 concludes.

2 Delays in Formal Standard Setting

SSOs are governed by formal and informal rules for reaching a consensus, which typically implies support from a substantial majority of participants. The American National Standards Institute (ANSI) offers the following definition:

> "...substantial agreement has been reached by directly and materially affected interests. This signifies the concurrence of more than a simple majority, but not necessarily unanimity. Consensus requires that all views and objections be considered, and that an effort be made toward their resolution." (ANSI, 2006)

Since SSOs are voluntary, the consensus principle gives interested parties the ability to block (or at least delay) the adoption of a new standard. This leads to lengthy delays when

[4] Many bargaining models provide alternative explanations for delay. For examples, see the review by Kennan & Wilson (1993) or the more recent contributions of Busch & Wen (1995) and Yildiz (2004).

participants have a private agenda, as is often the case. Some firms hope to profit by licensing their intellectual property. Others have made technology-specific investments and do not want to incur the costs of product redesign. More broadly, firms can benefit from learning economies, proprietary complements, or time-to-market advantages when their own technology becomes an industry standard. This section uses a simple model to analyze the link between private interests and coordination delays.

2.1 A Model of Consensus

Consider a symmetric committee with k members. Each sponsors a single proposal, and only one can be chosen as the industry standard. Proposals vary in quality, which is indexed by a random variable q with the symmetric joint distribution $F(\cdot)$ and continuous support. The element q_i is a publicly observed measure of player i's proposal quality. In practice, q_i will depend on product market characteristics and aspects of the underlying design; such as technical performance, implementation cost, and flexibility.

Committee decision-making is based on the stochastic bargaining model of Merlo & Wilson (1995). At the start of period t, committee members pay a participation cost c, and receive an independent draw $q_i(t)$ from the quality distribution. If $q_i(t)$ improves on i's last proposal, this draw becomes their new design (i.e. they are free to discard poor ideas). In the second half of each period, the player with the best proposal makes a take-it-or-leave-it offer, which may include concessions $b \geq 0$, and the remaining committee members vote whether to accept it. The consensus principle is modeled as a unanimity rule: standardization occurs when all players accept an offer.[5] If anyone votes to reject, the committee members receive an inside payoff — which I normalize to zero — and the game moves to another period, with all players applying a discount factor $\beta < 1$.

If player i's offer is accepted, they receive a payoff $\Pi_i(q) = \pi(q_i, a_i; \omega) - (k-1)b$ in the current period; where $a_i = 1$ indicates that i's proposal was chosen, and $\omega \geq 0$ is a parameter that measures the private benefits of winning. The remaining committee members get $\Pi_j(q) = \pi(q_i, a_j; \omega) + \gamma b$, where γ reflects the efficacy of any concessions.

To capture the "public goods" aspect of formal standardization, I assume that π is increasing in q; so everyone benefits from choosing a better standard. But participants also have a private interest in their own technology. Specifically, $\pi(q, a; \omega)$ is (weakly) increasing in a, and has increasing differences in (a, ω). I also assume that $\pi(q, 1) = \pi(q, 0)$ when $\omega = 0$. These are

[5] I model the bargaining process as an ultimatum-game where the player with the best proposal has all of the bargaining power. While engineers often highlight the importance of "technical merit" in SSO deliberations, a symmetric random-recognition rule would weaken the link between quality and bargaining power without altering the results.

weak assumptions that can accommodate many models of downstream competition, e.g. losers incur a fixed cost of redesign or pay royalties to the winner. The key point is that *ex post* payoffs become increasingly asymmetric as ω grows large.

As part of the bargaining process, committee members may use concessions to "buy off" opposition to their proposal. SSOs typically encourage certain types of concession, such as technical compromise or commitments to license intellectual property on liberal terms. However, these measures may not be efficient (or credible), and SSOs often limit the scope of bargaining — especially where monetary bribes are concerned.[6] When side-payments are costly, as I henceforth assume, then $\gamma < 1$. In practice, this parameter will depend on a variety of factors; notably technological trade-offs, antitrust policy towards SSOs, and the costs of licensing.[7]

2.2 Equilibrium

I analyze symmetric Markov-perfect equilibria, where each player's strategy is a correspondence σ_i that maps the state variable $q(t)$ onto an offer or voting rule. If $p \equiv p(q(t), \sigma)$ is the equilibrium probability of adopting a standard, these strategies must solve the following program:

$$V_{it}(\sigma) = -c + \max_{\sigma_i} \int_q \{p\Pi_i(q) + (1-p)\beta V_{it+1}(\sigma)\}\, dF(q) \tag{1}$$

Because $F(\cdot)$ is stationary and symmetric, we can write $V_{it} = V$. This implies a unique equilibrium at the voting stage: the committee will accept a proposal of quality x if and only if

$$\pi(x, 0) + \gamma b \geq \beta V \tag{2}$$

Given this voting rule, a proposer can capture any available surplus by offering concessions $b(x) = \max\{0, \frac{\beta V - \pi(x,0)}{\gamma}\}$, subject to the constraint that

$$\pi(x, 1) - (k-1)b(x) \geq \beta V \tag{3}$$

[6] Reluctance to allow monetary side-payments (or even explicit bargaining) is often linked to fears of antitrust litigation. Recent policy changes may increase γ by addressing these concerns. Specifically, the Standards Development Act of 2004 (H.R. 1086) gives registered SDOs immunity from triple damages in antitrust lawsuits; and the U.S. Department of Justice has issued Business Review Letters that grant explicit permission for prospective disclosure of patent licensing terms within an SSO.

[7] An alternative model that yields similar conclusions assumes that $\gamma = 1$ but players face a budget constraint, and can only make (or credibly commit to) concessions below some threshold $\bar{b}$. The expected costs of participation provide a natural upper-bound, since any more would presumably lead to an influx of participants who are primarily interested in collecting side-payments.

Since a committee that adopts x would also accept any better proposal, equations (2) and (3) define a reservation rule: the probability of achieving consensus is $p = 1[x \geq q^*]$, where q^* is the lowest-quality proposal that will be adopted in equilibrium. When $x = q^*$ then (3) must bind, or the proposer would be willing to offer additional concessions to secure approval, thereby lowering q^*.

Replacing the reservation rule p in the Bellman equation (1) yields V as a function of the equilibrium quality threshold. Plugging that expression into equations (2) and (3) and summing the system of k equalities leads to the following equilibrium condition, where $G(\cdot)$ is the cumulative distribution of x (i.e. the first-order statistic of F); $S(x) = \sum_i \pi(x, a_i)$ is the players' gross surplus; and $\Delta(x) \equiv \beta^{-1} - G(x)$:

$$kc = \int_{q^*}^{\infty} S(x)dG(x) - \Delta(q^*)S(q^*) + (\gamma - 1)(k - 1)\left\{\int_{q^*}^{\infty} b(x)dG(x) - \Delta(q^*)b(q^*)\right\} \quad (4)$$

The main results follow directly from equation (4). I state and discuss these predictions here, and provide formal proofs in Appendix A.

P1: There is a unique symmetric equilibrium.

P2: The equilibrium quality-threshold q^* exceeds a jointly optimal reservation rule q^{fb} if and only if there is distributional conflict ($\omega > 0$) and side-payments are costly ($\gamma < 1$). Intuitively, when $\omega > 0$ losers want to bargain beyond q^{fb}, since it may lead to the adoption of their preferred system (or equivalently, an increase in bargaining power). And when $\gamma < 1$, a proposer with quality q^* will not offer concessions that fully compensate for these opportunity costs, since that would lead to a payoff below the continuation value.

The proof of P2 is straightforward. When $\gamma = 1$, this is a bargaining game with transferable-utility. Since the last term in (4) equals zero, the expression is identical to the optimal stopping rule in a single-agent search problem (e.g. Lippman & McCall, 1976), so $q^* = q^{fb}$. The last term in equation (4) is also zero when $\omega = 0$, since side-payments are unnecessary when winning and losing payoffs are identical. But if $\omega > 0$ and $\gamma < 1$, the last term in (4) is strictly positive. Since the right side of equation (4) is decreasing in q^*, equality is restored at some $q^* > q^{fb}$.

P3: When ω is small, inefficient delays are (weakly) increasing in the level of distributional conflict. This is a corollary of P2: since $q^* > q^{fb}$ requires conflicting interests and costly side-payments, the two cutoffs must converge as ω approaches zero. (While extra assumptions on π will produce models where $G(q^*) - G(q^{fb})$ is always increasing with distributional conflict, it is not crucial here, since $\omega \approx 0$ is a reasonable characterization of the early IETF.)

P4: Equilibrium delays are (weakly) decreasing in the efficacy of side-payments. This result follows from totally differentiating equation (4) with respect to γ. The inequality is strict when

$\omega > 0$. Since γ does not influence q^{fb}, delays are more efficient (on average) when side-payments work better.

P5: As k grows large, the rent-seeking incentive disappears. Intuitively, a larger committee size reduces each player's odds of winning, so concessions decline relative to payoffs and direct costs. This implies that q^* will approach q^{fb}. However, the impact on expected delay is ambiguous, since adding players also changes $G(\cdot)$, the first-order statistic of the proposal quality distribution.

2.3 Discussion

The empirical analysis treats rapid Internet commercialization (and its impact on the IETF) as an exogenous shock to ω that can be used to test P3. The results can also be interpreted as a test of $\gamma < 1$, since there should be no correlation between conflict and delay if transferable utility leads to efficient outcomes. Before turning to the data, I pause to comment on several simplifying assumptions in the model.

First, since committee members have no outside option, failure to reach a consensus always leads to another round of bargaining. This would be the case, for example, if SSO participants are highly risk-averse and prefer a more competitive "open standards" environment to a winner-take-all standards battle. In practice, there may be a range of intermediate options, such as pushing for a parallel specification or seeking the endorsement of a different SSO.[8] While it would be interesting to model these scenarios, they do not pose problems below. Since failing to reach a consensus is observationally equivalent to $q^* = \infty$, I treat failed efforts as truncated observations.

Second, I analyze symmetric equilibria in an *ex ante* symmetric model that best describes committees with a few evenly matched players. Large *ex ante* asymmetries will presumably lead to a swift resolution in favor of the player with better technology, more resources, or stronger incentives to push their own solution, as in Myatt (2005). At some parameter values, the model will have asymmetric equilibria that favor a particular player. These outcomes suggest the alternative institution of a predetermined platform leader, such as IBM during the mainframe computing era. In practice, committees with many participants are often characterized by a small number of roughly symmetric coalitions. I leave the interesting and complex topic of coalition-building for future work, and simply add a variety of controls for firm-size to the empirical models.

Finally, I treat q as a non-contractible quality measure that reflects the expected profitability

[8] The scope for strategic forum shopping may be limited in the case of the IETF, since its early success helped to establish it as the *de facto* SSO for Internet standards.

of adopting a particular technology. Engineers often stress the importance of "technical merit" in formal standard setting. Moreover, contracting on quality will be difficult when it is hard to distinguish standards' underlying technical merit (i.e. selection effects) from the network effects unleashed by SSO endorsement. Unfortunately, the same measurement problem arises in empirical research. While I use citations to examine the *ex post* impact of IETF standards, the results can only provide suggestive evidence on the link between distributional conflict and the quality of outcomes, since it is hard to determine whether these cites are a cause of ω or a consequence of q.

3 The Internet Engineering Task Force

The IETF creates and maintains the standards used to run the Internet, which include TCP/IP (the Internet's core routing protocols); the Domain Name System (used to find other computers), and MIME (used to parse the content, formatting, and routing information contained in e-mail messages). This section provides a brief overview of the IETF and its standard setting process.[9]

3.1 History

The first IETF meeting took place in 1986, and was attended by 21 government funded researchers. Commercial interest remained limited for several years; partly because the National Science Foundation's Acceptable Use Policy prohibited commercial use of the network. By 1990, IETF meetings attracted 100 regular participants, with "... about 1/3 from vendors, about 1/3 from government (DoD and civilian agencies), and over 1/4 from universities and regional network operators" (IETF, 1990).

In the early 1990s, TCP/IP emerged as the *de facto* standard for computer networking, and the IETF evolved from a small quasi-academic networking community into a high-stakes forum for technical decision-making. This led to an increase in the number of committees and proposals (see Table 1), and to substantial changes in the demographics of IETF participants. Figure 1 graphs the the "suit-to-beard" ratio (i.e. the share of participants affiliated with private-sector organizations, as indicated by their email address), which increased from 55 percent in 1993 to 80 percent by late 2001.[10]

[9] Detailed histories of computer networking include Abbate (1999) and Berners-Lee & Fischetti (1999).

[10] See Appendix A for details on the construction of this variable, which may understate changes in IETF demographics, since it does not capture variation in the types of firm or employee (e.g. research scientist vs. consultant) within the Dot-com or Dot-net top-level domains. Table R-2 in the online appendix illustrates a similar trend in the institutional affiliation of proposal authors.

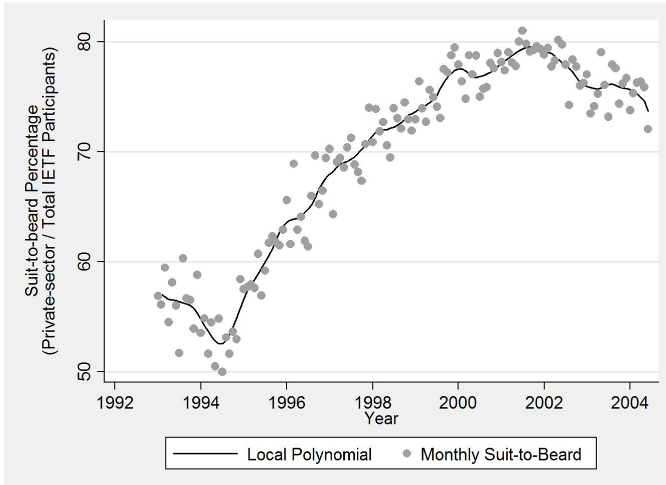

Figure 1: Commercial Participation in the IETF

Anecdotal evidence suggests that growth and commercialization led to increased tension within the IETF (e.g. Davies, 2004). For example, a committee working on instant messaging protocols around 2000 received proposals from both Microsoft and AOL, who were in the midst of a battle over the issue. Another committee working on protocols for calendar sharing applications witnessed significant disagreements among Microsoft, Lotus, and Netscape—each of whom was marketing proprietary scheduling software.

While the size, demographics and commercial significance of the IETF changed dramatically during the 1990s, its standard-setting process did not. Thus, this organization provides a unique opportunity to study the impact of distributional conflict on voluntary standard setting.

3.2 Standard Setting Process

Like most SSOs, the IETF has developed its own language for describing the standards process. Committees are called Working Groups (WGs) and proposals are called Internet Drafts (IDs). If an Internet Draft is approved, it gets published as a Request for Comments (RFC). Bradner (1996) contains a detailed description of IETF procedures. Figure 2 provides an overview of the process and associated terminology.

The IETF's unofficial motto is "rough consensus and running code." In the first stage of the standards process, participants identify a problem and form a Working Group to consider potential solutions.[11] Internet Drafts are circulated using email discussion lists and go through

[11] Individual authors can submit an unsolicited Internet Draft (and many do), but these are unlikely to become a RFC without Working Group approval (see Table R-3). The few that do are not considered "IETF standards."

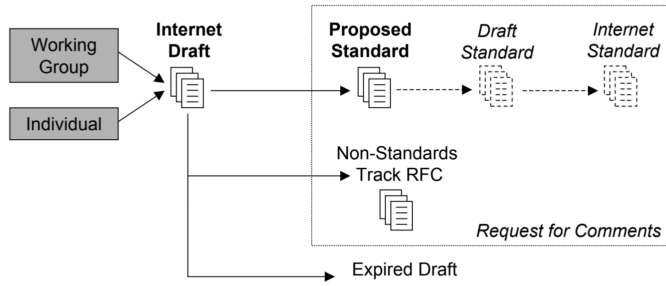

Figure 2: The IETF Standard Setting Process

a series of revisions. If the Working Group reaches a "rough consensus" on the merits of a proposal, the document is published as an RFC. In the case of a Proposed Standard, successful implementation and widespread deployment may provide sufficient evidence of "running code" to justify elevating a protocol to the status of Draft or Internet Standard.

Anyone can participate in the standards development process by attending meetings or joining the debate on a Working Group's e-mail discussion list. However, the final decision to publish an RFC rests with the WG chair and an IETF-wide advisory board called the Internet Engineering Steering Group (IESG).[12] The chair's job is to determine whether the committee has reached a consensus. There is no formal voting. Rather, the chair issues a "last call" for comments from WG members and submits a proposal to the IESG. The IESG reviews the Internet Draft and issues a "last call" for comments from the entire IETF community. Any comments or formal appeals go to the IESG, which decides whether to publish. The review process emphasizes dispute resolution within the Working Group, and discussions with a number of participants suggest that RFC publication implies support from a substantial majority of the relevant committee.

To emphasize consensus decision-making, I focus on the first stage of this process, which begins with the submission of an Internet Draft and culminates in one of three ways: publication as a Proposed Standard, publication as a nonstandards-track RFC, or expiration.

3.3 Standards and Nonstandards

The key distinction between Proposed Standards and nonstandards-track RFCs is that only standards receive the IETF's formal endorsement. This makes them more commercially relevant — and more likely to generate distributional conflict — than nonstandards.

[12]The IESG has roughly twenty members, including the IETF chairperson, six Technical Area directors, and several liaison and ex-officio members.

Proposed Standards are meant to be prescriptive; a signal to technology developers that, "if you are going to do something like this, you must do exactly this" (Postel, 1995, pg. 10). They are also an explicit call for IETF members to produce "running code" by incorporating the new standard into their own products and services. To assess compliance with a Proposed Standard, the IETF uses a set of "standards-track key words" to define the formal requirement-level for each feature defined in a new protocol.[13]

Nonstandards-track RFCs are an outlet for publishing new ideas without formally endorsing them. The IETF divides nonstandards into two categories: Informational and Experimental (though in practice, the distinction is rather vague).[14] These purely informative documents cannot advance to the status of Draft or Internet Standard without first becoming a Proposed Standard. And the IETF does not provide any mechanism for assessing "compliance" with nonstandards, since they merely describe one possible way of doing things.

Working Groups typically use the nonstandards-track in two ways. The first is to publish ideas that are too preliminary or controversial to become a Proposed Standard. For example, one Experimental RFC that eventually became a Proposed Standard (RFC 2582) suggested changes to the Transmission Control Protocol (TCP) to help manage network congestion. While the IETF did not initially endorse this idea, it was published as a nonstandard to encourage further experimentation. A few ideas default to the nonstandards-track because they cannot achieve consensus. One extreme example is the Informational "Link-Local Multicast Name Resolution Protocol" (RFC 4795), which went through 48 revisions as an Internet Draft.

A second common use of nonstandards is to provide information that complements a standard, such as guidelines for implementation and deployment. For instance, an Informational RFC entitled "Known TCP Implementation Problems" (RFC 2525) catalogs various negative externalities that occur when vendors fail to comply with different parts of TCP. Another highly-cited Informational RFC describes "An Architecture for Differentiated Services" (RFC 2474) based on protocols defined in a set of related standards.

While I cannot observe whether a new proposal is meant to become a Proposed Standard or nonstandards-track RFC, discussions with IETF participants suggest that most authors have a clear idea after one or two revisions (if not sooner). Moreover, there can be no ambiguity about an ID's intended publication-status in the final stages of the review process, since that information is provided in each of its "last call" announcements.

[13]The key-words are "must, must not, required, shall, shall not, should, should not, recommended, may, and optional" in all capital letters (Bradner, 1997).

[14]A third type of nonstandard-track RFC, called a Best Current Practice, is used to document IETF procedures. I exclude these from the analysis.

4 Empirical Strategy

This section describes a strategy for identifying the link between commercialization and delays in the IETF standard-setting process, and discusses the data used in estimation.

4.1 Identification and Estimation

Suppose we observe data from a cross-section of proposals i, submitted to committees j at times t. These data contain the time-to-consensus T_i^s (where s indicates a standard) and a vector of committee and proposal characteristics X_{ijt} that includes measures of distributional conflict. To estimate the impact of rent-seeking on delays, one might use a simple linear model:

$$T_i^s = X_{ijt}\beta^s + \varepsilon_i^s \tag{5}$$

This approach is straightforward, but vulnerable to omitted variables. The direction of any bias is not clear *a priori*. For example, if X is positively correlated with both distributional conflict and unobserved technological complexity, β^s is likely to overstate the rent-seeking effect. On the other hand, if X is positively correlated with both conflict and a sense of urgency, then β^s could be downward biased.

I address this endogeneity problem by using nonstandards-track RFCs to estimate a model of counterfactual durations. Suppose T_i^n represents the duration of proposal i in the absence of distributional conflict. The rent-seeking effect is the derivative of $E[T^s - T^n|X]$ with respect to X. Letting ε_i^n represent unobserved factors that lead to routine delays, we have:

$$T_i^n = X_{ijt}\beta^n + \varepsilon_i^n \tag{6}$$

If S_i is an indicator variable that equals one for Proposed Standards and zero for nonstandards-track RFCs, then T_i can be written in terms of (5) and (6) as:

$$\begin{aligned} T_i &= S_i * T_i^s + (1 - S_i) * T_i^n \\ &= X_i\beta^n + X_iS_i\delta + \eta_i \end{aligned} \tag{7}$$

When X contains only a constant and time-dummies, (7) is the familiar difference-in-differences model. When X includes continuous proxies for distributional conflict, the rent seeking effect $\delta \equiv \beta^s - \beta^n$ measures a difference in slopes rather than means. But the intuition is unchanged: if $X_i\beta^n$ measures the relevant counterfactual (i.e. routine delays), and S is uncorrelated with η, then δ identifies the impact of distributional conflict on coordination delays.

In a controlled setting, one might implement (7) by randomly selecting Proposed Standards and devising a treatment that removes any rent-seeking incentives without altering the technical problem or publication process. As a practical alternative, I use nonstandards-track RFCs to construct a "no conflict" control sample. Since nonstandards and standards describe the same technology and go through an identical publication process — often within the same Working Group — they should be subject to the same routine delays. However, the commercial stakes are much lower for nonstandards-track RFCs, since they are not meant to provide an impetus for product-market coordination. I therefore assume that nonstandards-track delays provide an observable measure of T_i^n that allows me to estimate (7).

The major concern with this approach is that Internet Drafts are not randomly assigned to a particular publication-type. If the certification benefits of becoming a Proposed Standard are constant, while delays vary with the amount of distributional conflict, then selection should produce downward-biased estimates of δ. Intuitively, authors who anticipate an especially rancorous debate will steer their idea onto the nonstandards-track; forgoing certification benefits in favor of a less costly publication process.[15] While the nonstandards-track does provide an outlet for ideas that fail to achieve consensus, one could imagine that certification benefits are highly correlated with distributional conflict, leading to a bias in the opposite direction. I use several methods to guard against this possibility.

First, I use propensity-score matching (Rosenbaum & Rubin, 1983) to control for selection on observables. Second, I include Working Group fixed-effects to control for time-invariant committee-level unobserved heterogeneity (e.g. the complexity of the underlying technology). And finally, I estimate an endogenous switching regressions model that uses standards-track keywords as an instrument for S. The switching model begins with the two duration equations (5) and (6), and adds a selection equation:

$$S_i = 1[Z_i\pi + \nu > 0] \tag{8}$$

I assume the error-terms have a joint-normal distribution: $(\varepsilon^s, \varepsilon^n, \nu) \sim N(0, \Sigma)$, where σ_ν^2 is normalized to one; the covariances $\sigma_{\nu s}$ and $\sigma_{\nu n}$ are parameters to be estimated; and the covariance between the two duration equations is undefined.[16] I also assume that every proposal would eventually become an RFC, so the probability of observing a censored or expired proposal is $Pr(S_i = 1, T_i^s > T) + Pr(S_i = 0, T_i^n > T)$. The switching model therefore accomplishes two

[15] Similarly, if participants with more controversial proposals try to form coalitions in advance of submitting a draft, this should lead to shorter standards-track durations.

[16] This model was first implemented by Lee (1978) and is discussed in (Maddala, 1983, pg. 223). The online appendix provides additional details on the likelihood function and estimation.

goals: correcting for sample selection (as in Heckman, 1979), and incorporating expired and censored proposals, which I omit when estimating (7) because their intended publication status (S_i) is never observed.

Standards-track keywords play an important role in the switching model, since I assume this variable enters (8), but not the duration equations. Keywords thus act like an instrumental variable; they influence delay only through their effect on S_i. Intuitively, keywords provide a measure of the authors' intentions, and the exclusion restriction will be valid if the "idea-generating process" that determines standards-track suitability is orthogonal to unobserved factors that influence delays. If keyword frequency is positively correlated with both both S and $\varepsilon^s - \varepsilon^n$ (e.g. because keywords reflect the scope or specificity of a proposal, which is not relevant for nonstandards but results in longer standards-track delays), the switching model could produce upward-biased estimates of the rent-seeking effect. In practice, I find the correlation between keywords and delay to be slightly higher for nonstandards than standards, though neither relationship is statistically significant. This suggests that keywords are a reasonable instrument for S, as well as a practical solution to the problem of modeling censored and expired proposals.

4.2 Data and Measures

All of the data come from the IETF's public archives. The population consists of 3,521 Internet Drafts submitted to IETF Working Groups between January 1993 and December 2003. I restrict attention to an estimation sample of 2,601 IDs that went through at least one revision.[17] While there are 249 Working Groups in the estimation sample, 25 of them fail to publish any RFCs, and only 176 publish more than one. The median Working Group evaluated seven proposals, and published one Proposed Standard and one nonstandards-track RFC. The largest Working Group (IP Security) considered 123 proposals and published 54 RFCs.

The IETF's file-naming conventions identify the Working Group and version number of each proposal. For example, the file "draft-ietf-mmusic-rtsp-06.txt" corresponds to the sixth revision of the Real-Time Streaming Protocol (rtsp) produced by the Working Group on Multiparty Multimedia Session Control (mmusic). Submission and revision dates for each ID were obtained from the "ietf-announce" listserv, which is used to announce all new IETF publications. I track proposals in the estimation sample between January 1993 and June 2008, and the primary dependent variable (Total Days) measures the time between initial submission and final revision.

[17] I exclude 61 IDs submitted to "non-technical" Working Groups (from the IETF's General and User Services Areas) and 859 drafts that were never revised. While I cannot calculate a duration for unrevised proposals — of which, only 79 were published — including them in the regressions with an arbitrary (but short) duration does not change any results. See Table R-3 in the online appendix for additional details on sampling.

Table 1 shows mean publication lags by publication-type and submission-year. The average time-to-consensus for a Proposed Standards was 774 days (2.1 years), compared to 595 days for a nonstandards-track RFC and 487 days for an expired proposal. While delay is clearly increasing on both tracks, a simple OLS regression shows a statistically significant difference in time-trends; with average duration growing by 49 days per-year for standards and 31 days per-year for nonstandards. Table 1 also shows that Proposed Standards submitted in 1997 and 1998 experience a sharp drop in publication lags relative to nonstandards. This may reflect the urgency surrounding standards that were crucial to large network operators' expansion plans, or the "fast tracking" of several web-related protocols.[18] However, I estimate models that allow the rent-seeking effect δ to vary by Draft-Cohort and find no evidence of a systematic decline during this period.

Table 2 provides a short definition and summary statistics for all variables used in the analysis. (Appendix B provides additional details on data collection.) The top panel contains a variety of outcome measures, while the second panel presents two measures of distributional conflict based on IETF commercialization. Note that both measures of distributional conflict vary at the Working Group level, which should alleviate concerns that they are endogenous to proposal-level factors that influence delay.

The first measure of distributional conflict is a WG-level "suit-to-beard ratio" (see Figure 1). Suit-share is defined as the percentage of all email domains (e.g. ibm.com or berkeley.edu) from dot-com or dot-net organizations on a WG listserv during a one-year window prior to the initial submission of an ID.[19] The idea behind this measure is that standard-setting creates more distributional conflict as the underlying technology gets closer to commercial application, since firms must commit to a particular design and the standard becomes more likely to impact existing products. Thus, increased participation by organizations in the dot-com and dot-net top-level domains is a good indication that a Working Group is producing commercially relevant technology.

While I interpret Suit-share as a measure of distributional conflict, it may also capture variation in the size of the market. The model (specifically P2) predicts that the magnitude of payoffs will only influence (inefficient) delays when the benefits of choosing a standard are asymmetric. Moreover, larger payoffs will produce faster decisions in a model without conflict, since they increase the opportunity cost of delay. In practice, increasing market size is likely to produce more distributional conflict when firms recognize the opportunity and race to reach

[18] These hypotheses were suggested in an private email communication from a former IESG member. Examples include the RADIUS protocol for dial-in authentication, and SSL for web security.

[19] For international participants, I used a series of country-specific rules to classify their top-level domain. For a discussion of spam, hosted mail, and related issues, see Appendix B.

the market first with a proprietary solution.

As a second proxy for distributional conflict, I create a measure of "background IPR" by linking each ID author's email address to an assignee code in the NBER U.S. patent database. For each assignee, I calculate a five-year cumulative patent stock, and weight that by the uncentered correlation between the assignee's patent portfolio and the cumulative patent portfolio of all IETF participants (based on three-digit USPTO technology classifications).[20] To create the WG-level log(Patents) variable, I sum this weighted average patent stock over all firms with one or more proposals before a Working Group in a given year.

The third panel in Table 2 contains three dummy variables that indicate whether an Internet Draft had one or more authors from the dot-org, dot-edu, or dot-gov top-level domains. These authors are typically members of the academic, government and non-profit research communities, and I use the indicator variables as a proxy for the efficacy of side-payments. This interpretation is based on the idea that non-commercial participants — whose efforts are typically limited to one or two standards — will find it difficult to offer concessions that appeal to large technology vendors, since they cannot use cross-licensing or log-rolling to craft a compromise that spans multiple markets or committees.

The last panel in Table 2 summarizes a number of additional control variables. These include a Draft-Cohort variable that measures an ID's initial submission year; variables that capture WG activity and prior experience; measures of proposal size and complexity; dummies for several author-attributes (notably whether any author served as a WG chair); and dummies for the six Technology Areas defined by the IETF.

4.3 Comparing Means

The "difference in slopes" model of equation (7) assumes that assignment to the standards-track is exogenous. If this is true, the sample means of all predetermined variables should be the same for Proposed Standards and nonstandards-track RFC's. While these means are indeed quite similar, I use propensity-score matching to create a sample where they are balanced by construction.[21]

Table 3 compares sample means for Proposed Standards and nonstandards-track RFCs in the full and matched samples. The top panel examines outcomes (which should not necessarily balance). Proposed Standards take more time to publish; go through more revisions; are

[20] The weighting places more emphasis on the patents of firms that are close to a hypothetical "IETF average" technology profile, and reduces the influence of outliers. I also removed some the largest firms, such as IBM, and found that it did not make a difference.

[21] Specifically, I drop RFCs where the probability of becoming a Proposed Standard (based on fitted values from a probit) is below the 10th percentile of the empirical distribution for Proposed Standards, or above the 90th percentile for nonstandards-track RFCs. See the online appendix for additional details.

mentioned in more emails; and contain more key words than nonstandards-track RFCs. These differences do not disappear in the matched sample. Proposed Standards also receive more forward-citations from RFCs and U.S. patents, which is broadly consistent with the assumption that they have greater commercial impact.[22] Interestingly, standards and nonstandards are cited at roughly the same rate by academic journal articles. This may reflect the fact that academic authors are more likely to produce nonstandards-track RFCs, or perhaps they are less concerned with the commercial implications of the underlying idea.

The bottom half of Table 3 shows that for most of the control variables, a T-test does not reject the hypothesis of equal means. Of the five main proxies for distributional conflict or access to side-payments, the only meaningful difference appears on Dot-edu, indicating that academics are more likely to produce nonstandards. The full sample T-tests also show that Proposed Standards are produced by somewhat "older" and "busier" Working Groups (i.e. committees that have evaluated more Internet Drafts and generated more email traffic). While these differences are not particularly large, they are statistically significant, and Hotelling's omnibus test rejects the hypothesis that the sample means of all variables are equal. For the matched sample, however, Proposed Standards are statistically indistinguishable from nonstandards-track RFCs. Thus, Table 3 suggests that differences in the sample space across standards and nonstandards are unlikely to drive any difference in estimated coefficients, particularly for the matched sample.

5 Results

This section presents evidence that Internet commercialization led to increased distributional conflict and slower standards production at the IETF. There are three main findings. First, commercialization (as measured by Suit-share, and to a lesser extent, background IPR) is positively correlated with an increasing difference between standards and nonstandards-track delays. Second, this rent-seeking effect increases over time, and is larger for standards that are more likely to interface with proprietary technologies. And third, delays are correlated with an increase in citations to Proposed Standards, but not nonstandards-track RFCs.

5.1 Distributional Conflict and Coordination Delays

Figure 3 provides a simple non-parametric illustration of the main result: as Suit-share increases from 60 to 100 percent, the average standards-track publication lag increases steadily, while

[22] While Proposed Standards are cited at nearly twice the rate of nonstandards, the difference increases to roughly 400 percent for Draft and Internet Standards.

nonstandards-track delays exhibit little or no change. To control for time-trends, technology effects or other factors that might influence this pattern, I turn to a parametric model based on equation (7). As described above, this difference-in-slopes estimator assumes that selection onto the standards-track is exogenous, and estimates the rent seeking effect by interacting measures of distributional conflict with a standards-track dummy variable. The results are presented in Table 4.

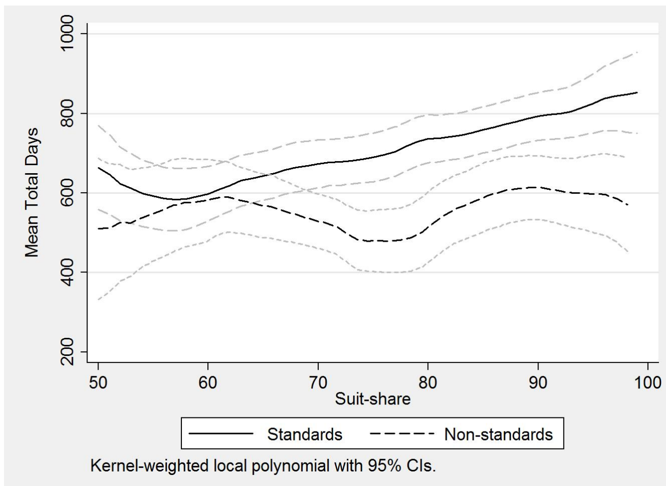

Figure 3: Difference in Duration by Suit-Share

Model (1) shows OLS estimates for all RFCs submitted between 1993 and 2002 with a publication lag less than 5.5 years (to control for right-censoring of the dependent variable). Since this specification includes proposal-year and technology-class effects, but not WG fixed effects, the parameters measure both within and between committee variation in distributional conflict. The interaction of Suit-share with an S-Track indicator is large and statistically significant. In particular, a one percentage-point increase in Suit-share adds 8.3 days to the publication process. The log(Patents) interaction is smaller; doubling the level of background IPR increases delay by roughly 12 days. While the latter effect is not significant, this is due to imprecision in the nonstandards-track baseline: restricting the main effect of log(Patents) to zero causes the interaction to become significant at the 5 percent level.

Model (1) also includes a set of author attributes (Dot-gov, Dot-org, and Dot-edu) that proxy for access to side-payments. Having a Dot-org or Dot-edu affiliated author on a Proposed Standard adds between 5 and 7 months to the publication process (relative to the nonstandards-track baseline). This result is consistent with P4, which predicts longer delays as the efficacy of concessions declines.[23]

Model (2) uses the same specification, but restricts attention to the matched sample and weights each observation by the inverse of its propensity score. All of the parameter estimates are statistically indistinguishable from model (1). Model (3) continues to use the matched sample, but drops the propensity-score weights and adds Working Group fixed-effects to control for time-invariant committee-level heterogeneity.[24] Once again, the estimates change very little. While an F-test clearly rejects the null-hypothesis that the fixed-effects are jointly equal to zero, a Hausman test indicates that the OLS estimates in (1) are consistent.

The last two columns in Table 4 drop the nonstandards-track baseline and return to the full sample of Proposed Standards. The log(Patents) coefficient increases in both size and significance, while the Suit-share, Dot-org and Dot-gov effects decline slightly. The results do not change if I add WG fixed-effects. Comparing (4) and (5) to the diff-in-diffs models suggests that IETF commercialization led to slower standards production, but slightly faster publication of nonstandards-track RFCs. One interpretation of the latter finding is that a faster nonstandards-track publication process reflects increased urgency created by the rapidly increasing scale of the overall network.

All of the results in Table 4 are robust to a wide variety of changes in specification and measurement. In particular, one can add large firm fixed-effects; change the dependent variable to Versions (i.e. a count of revisions); or estimate hazard models on a sample that includes the right-censored RFCs with especially long publication lags. Using one-year lags of Suit-share as an instrumental variable in model (4) produces slightly larger estimates of the rent-seeking effect. Replacing the WG-level Suit-share variable with ID Suit-share (a measure based on emails that mention a specific proposal while it is under evaluation) generates similar estimates of the rent-seeking effect. Finally, if one aggregates these data to the WG-Year-Track level and estimates a panel model with WG fixed-effects, similar to (3), the Suit-share interaction remains statistically significant.[25]

5.1.1 Switching Model

Table 5 presents maximum-likelihood estimates of the switching model described in Section 4.1. This model relaxes the assumption of exogenous assignment to the standards-track, and incorporates censored and expired proposals. Panel (1) reports the switching model parameter

[23]While I restrict the nonstandards-track coefficients on Dot-org, Dot-gov and Dot-edu to be equal, the results are quite similar if I allow the main effects for each of these indicator variables to vary.

[24]Models that include Working Group by publication-type interaction effects produce similar coefficients, but very large standard errors, since most Working Groups only publish a handful of RFCs (see Table R-4 in the online appendix). In particular, very few committees produce a large number of nonstandards.

[25]All of these results are presented in Appendix Tables R-6 and R-5.

estimates, while (2) and (3) present a pair of OLS regressions for comparison.

The key assumption in (1) is that log(Keywords) captures variation in the intended publication-status of an Internet Draft, but can be excluded from the two duration equations. The data support this assertion. In particular, while log(Keywords) is highly significant in the standards-track selection equation, it has no effect in the two OLS duration models.[26]

Controlling for endogenous selection onto the standards-track does not change the Suit-share result. In particular, a Wald test strongly rejects the hypothesis that the standards and nonstandards-track coefficients in panel (1) are equal. However, coefficients on the author-level proxies for costly side-payments (Dot-org and Dot-edu), are no longer statistically different across tracks. Comparing panel (1) to models (2) and (3) shows a persistent, though imprecisely estimated, difference between the two Dot-org coefficients. The Dot-edu difference appears to to be almost entirely a selection effect.[27]

Since the switching model produces estimates of the correlation between unobserved error-terms across equations, it is possible to test the null hypothesis of an exogenous selection process, i.e. $\sigma_{s\nu}$ and $\sigma_{n\nu}$ are jointly equal to zero. While that test is rejected in Table 5, I find that the correlation parameters tend towards zero as the specification becomes more parsimonious. For example, I cannot reject exogenous selection in a specification where the 30 Draft-Cohort dummies are replaced by a set of second-order polynomials. Given these results — and the fact that controlling for selection does not alter the Suit-share effect — I return to the simpler OLS framework in the next subsection, which examines heterogeneity in the link between conflict and coordination delays.

5.2 Time and Technology Interaction Effects

To examine variation in the impact of distributional conflict on delays, I interact the rent-seeking effect (δ) with a set of time and technology dummies. This leads to the following specification

$$T_i = X_i\beta^n + X_i S_i \delta_t + \eta_i \tag{9}$$

where t indexes either Draft Cohorts or IETF Technology Areas. Given the previous results, I focus on the Suit-share measure of distributional conflict, rather than log(Patents). The results

[26] In a probit model of the selection process that yields estimates very similar to (1), a one unit change in log(Keywords) corresponds to a 6 percent increase in the probability of becoming a Proposed Standard (holding other variables at their sample means).

[27] Another possible explanation for changes in the Dot-org and Dot-edu coefficients is changes to the estimation sample. However, estimating the switching model on a sample that does not contain censored or expired proposals generates results very similar to panel (1).

are presented in Table 6.

Models (1) and (2) estimate a separate rent-seeking effect for each of the IETF's six Technology Areas. The coefficients are arranged according to the TCP/IP "protocol stack" — a conceptual model of the hierarchical layers in a communications network. In this model, Application protocols communicate with a Transport layer, which calls on the Internet layer, which interacts with the Routing layer. (Security and Operations protocols provide services that are shared by the entire stack.) While estimates of the rent-seeking effect vary from 3.4 to 9.7 days per Suit-share point across the six IETF technology areas, they are largest near the "edges" of the protocol stack (i.e. on either side of the Internet Area), and in the Operations Area, which develops tools for network management.[28] Model (2) relaxes the assumption that the nonstandards-track Suit-share coefficient is equal across technology classes, and finds essentially the same results.

The finding that rent-seeking is less pronounced in the center of the protocol stack is broadly consistent with (and perhaps a consequence of) the "End-to-End" principle described by Saltzer *et al.* (1984). This design principle states that additional functionality should always be added near a network's end-points, as opposed to its core; and is one of the key features that make the Internet a radical departure from centrally switched phone networks. End-to-end design and the emergence of TCP/IP as a *de facto* open standard encouraged widespread experimentation at the upper and lower boundaries of the stack (e.g. IP now runs over phone-lines, copper cable, fiber-optics, Wi-Fi, cellular and satellite). This experimentation can lead to conflict. For example, Routing Area protocols often share an interface with the proprietary solutions of router and switch vendors, such as Cisco, Nortel or Juniper, whose products are typically implemented in hardware, which often means longer lead-times and substantial (sunk) engineering costs.

Models (3) and (4) examine variation in the rent-seeking effect over time. Specifically, I group proposals into a series of two-year windows based on submission year (Draft-cohort) and estimate a coefficient on the three-way interaction between the standards-track indicator, Suit-share and a dummy for each time window. I drop the Working Group effects in these models, since they are are very demanding on the data — few committees experience large demographic changes over a short time period *and* initiate multiple standards and nonstandards during that window.[29]

The results show that rent-seeking has increased over time. Model (3) suggests a jump

[28] F-tests reject the null hypothesis that the Internet Area coefficient is equal to the Routing, Transport or Operations Area parameter at the 5 percent level.

[29] Including WG fixed effects produces qualitatively similar results that are not statistically significant due to very large standard errors.

around 1995 and again in 1999; while model (4), which allows the nonstandards-track Suit-share parameter to vary by cohort, exhibits a more steady increase. Neither model shows a sharp drop in the Suit-share interaction term for the 97-98 Draft Cohort, which is somewhat surprising given the summary statistics in Table 1. This suggests that the standards-track "dip" is primarily a compositional effect. In particular, web extensions may have led to an increase in relatively young, and perhaps less commercial, technical committees during that period; which could reduce standards-track publication lags without altering the relationship between Suit-share and coordination delays.

5.3 Coordination Delays and *Ex Post* Impact

The model in Section 2 predicts that distributional conflict is linked to both delay and the quality of outcomes through a more stringent selection process. This sub-section tests that hypothesis using forward-citations from RFCs, U.S. patents and academic journal articles as a proxy for the "impact" of IETF standards. I do not offer a specific interpretation for these cites, which may reflect knowledge flows, economic value, technical novelty, or other factors that vary according to the source of the citation.[30] However, I interpret a broad increase in citation-rates across patents, standards and journal articles as evidence that an RFC is more important, or "higher quality" in the broad sense implied by the model.

The citation analysis continues to use the diff-in-diffs framework of equation (7). However, since citations are highly skewed, I estimate Poisson regressions with conditional fixed effects at the technology-class by publication-track level. To control for variation in the significance of narrowly defined technologies, I introduce two covariates: log(ID Mail) and $\log(\text{Cites/Month})_{-i}$. The first variable measures the number of emails that mention a specific ID, and the second measures the average citation rate of other RFCs produced by the same committee (much like a WG fixed effect).[31] Since older RFCs naturally receive more citations, I include a fourth-order polynomial in months since publication, as well as a full set of Draft-cohort effects. The estimation sample includes all RFCs from Draft Cohorts 1993 through 2003, and results are presented in Table 7.

Models (1) through (3) use the sum of RFC, patent and journal citations as the dependent variable. While I find no relationship between Suit-share and aggregate cites, models (2)

[30] While there is a literature that examines the economic value of patent cites (e.g. Hall *et al.*, 2005), I focus on citations to non-patent prior art, which have received less attention. I do not know of any efforts to calibrate the economic value of citations from journal articles or compatibility standards.

[31] This specification is a compromise. It is hard to interpret the interaction terms in a Poisson model if the Standards-track dummy is not absorbed by the fixed effects. However, including WG by publication-track effects throws out much of the data, since many groups produce a single RFC (or receive no citations) on a given track. Estimating OLS models with WG fixed effects and an S-track dummy produces qualitatively similar results.

and (3) show that two measures of coordination delay — Total Days and Versions — are positively correlated with the difference between standards and nonstandards-track citation rates. Models (4) through (6) separate the three citation-types. I focus on Versions as a measure of delays, since the revision process is not necessarily co-linear with submission and publication dates (though the raw correlation with Total Days is 0.75).[32] In all three models, the relative citation rate of Proposed Standards increases with delay. However, there is no apparent relationship between Suit-share and standards-track citations. For nonstandards, commercialization is weakly correlated with RFC and patent cites, but leads to a drop in journal cites.

The Suit-share results are not consistent with P3 if we interpret cites as a measure of *ex ante* "technical quality." However, they are not surprising if one believes that standards produced in the early (less commercial) days of the IETF achieved a lasting impact partly because of their open-ness.[33] The relationship between delay and citations is consistent with P3; but there are two possible interpretations. First, unobserved distributional conflict may lead to a higher quality threshold, which causes both longer delay and more cites. A second possibility is reverse causation: cites measure conflict which produces delays. Two pieces of evidence favor the latter view. One is the insignificant Suit-share coefficient, which suggests a weak link, if any, between conflict and *ex ante* quality. And second, the strong correlation between *ex ante* email and *ex post* citation rates indicates that IETF participants have some idea of a proposal's likely impact before it gets published. Either way, these results provide some additional evidence that distributional conflict is linked to inefficient delays, since cites are only correlated with publication-lags on the standards-track.

5.4 Discussion

Overall, the empirical results suggest that the IETF's evolution from a quasi-academic institution into a high-stakes forum for technical decision-making led to increased "politics" and a slowdown of consensus decision-making. To calculate the magnitude of this effect, consider that Figure 1 shows a 25 percentage point increase in Suit-share. If each one-point increase led to 8 days of deliberation, then commercialization added roughly 200 days to the length of the IETF standard setting process. This effect would be economically significant in the market for networking hardware, where product life-cycles are often very short. However, given the remarkable growth of the Internet during this time period, one might reasonably conclude that

[32] While the main effect in model (2) is identified by functional form, since Total Days = Publication Date - Submission Date, the interaction with S-track would still be identified in a linear model.

[33] Furman & Stern (2006) provide evidence of an open-ness effect on citations in the context of biological research.

the IETF successfully managed a very challenging transition.

While these results clearly link distributional conflict to SSO performance, they do not have straightforward welfare implications. For one thing, they are silent on the costs or benefits of timely standard setting for end-users.[34] Moreover, credible welfare analysis would require a broader accounting of the costs and benefits of alternative institutional arrangements, such as a standards war or platform leader. Nevertheless, the consensus view of many in the standard setting community is that rapid technological change has increased the value of speed in cooperative standards development. These results suggest that the desire for speed may be incompatible with other objectives, such as open participation, consensus decision making, and the ability to retain a proprietary interest in proposed technologies.

6 Conclusions

This paper begins with a simple model of standard setting committees that emphasizes the trade-off between coordinating on a superior technology and owning a piece of the industry standard. The model shows that when concessions are costly and new standards create winners and losers, SSO committees will be hampered by the rent-seeking behavior of individual participants. This leads to a slower standard setting process, but may also produce higher quality outcomes.

I test these predictions using data from the IETF, and find evidence of strategic maneuvering within standard setting committees. In particular, distributional conflicts — driven by Internet commercialization — lead to a slower standard setting process. This link is strongest in areas where IETF standards are most likely to share an interface with proprietary technology, i.e. at the edges of the protocol stack. I also find that delays are correlated with the *ex post* significance of standards. However, it is unclear whether this is due to a more stringent selection process (as predicted by the model), or a link between the anticipated importance of a new technology and greater distributional conflict.

The results suggest that SSOs perform better in an environment that is free of the pressures created by imminent commercialization. They also show that SSO endorsements, while non-binding, matter enough for firms to fight over them. Thus, as shared technology platforms continue to gain commercial significance, we should expect to see more disputes over compatibility standards and individual firms' rights in the underlying technology.

The IETF's story also has broader implications for understanding the institutions that promote sharing and coordination in basic science and engineering. In particular, while commercial

[34] For example, delays that lead to a smoother migration path for end-users who are already part of some installed base could be efficient.

pressures have led to a slower standards process, the IETF continues to produce significant technology, and might even be viewed as a case study in effective governance of a shared platform. I suspect this success owes much to a unique culture that allows ideas to emerge from anywhere; places practical experience (i.e. "running code") ahead of top-down design; and uses mechanisms like the dual-track publication process to manage the tension between picking winners and encouraging widespread experimentation with new ideas.

Finally, this paper highlights several avenues for future research. One important question is 'Why do firms participate in SSOs?' While some firms clearly hope to influence outcomes, others may find that SSOs offer a convenient forum for collaborative R&D, a place to seek out new technical opportunities, or a chance to assess competitors. Another possibility is to study the relationship between SSOs' internal organization and standard setting outcomes. Finally, we have much to learn about the interaction between market and non-market standard setting — particularly the factors that determine whether a new standard emerges from an SSO, a *de facto* technology adoption process, or both.

References

Abbate, Janet. 1999. *Inventing the Internet*. Cambridge, Mass: MIT Press.

ANSI. 2006. *ANSI Essential Requirements: Due process requirements for American National Standards*. American National Standards Institute.

Becker, Gary S., & Murphy, Kevin M. 1992. The Division-of-Labor, Coordination Costs and Knowledge. *Quarterly Journal of Economics*, **107**(4), 1137–1160.

Bekkers, Rudi, Duysters, Geert, & Verspagen, Bart. 2002. Intellectual Property Rights, Strategic Technology Agreements and Market Structure: The case of GSM. *Research Policy*, **31**, 1141–1161.

Berners-Lee, Tim, & Fischetti, Mark. 1999. *Weaving the Web : The original design and ultimate destiny of the World Wide Web by its inventor*. 1st edn. San Francisco: HarperSanFrancisco.

Besen, Stanely M., & Farrell, Joseph. 1994. Choosing How to Compete - Strategies and Tactics in Standardization. *Journal of Economic Perspectives*, **8**(2), 117–131.

Besen, Stanely M., & Saloner, Garth. 1989. The Economics of Telecommunications Standards. *Pages 177–220 of:* Crandall, R., & Flamm, K. (eds), *Changing the Rules: Technological Change, International Competition, and Regulation in Telecommunications*. Washington: Brookings.

Besen, Stanley M., & Johnson, Leland. 1988. Compatibility Standards, Competition and Innovation in the Broadcasting Industry. *RAND Study R-3453-NSF*.

Bradner, Scott. 1996 (October). *RFC 2026: The Internet Standards Process – Revision 3*. Tech. rept. Internet Engineering Task Force.

Bradner, Scott. 1997. *Key words for use in RFCs to Indicate Requirement Levels*. Tech. rept. Internet Engineering Task Force.

Bresnahan, Timothy F., & Greenstein, Shane. 1999. Technological competition and the structure of the computer industry. *Journal of Industrial Economics*, **47**(1), 1–40.

Busch, Lutz-Alexander, & Wen, Quan. 1995. Perfect Equilibria in Negotiation Model. *Econometrica*, **63**(3), 545–65.

Cargill, Carl. 2001. Evolutionary Pressures in Standardization : Considerations on ANSIs National Standards Strategy. *Testimony Before the U.S. House of Representatives Science Committee*.

Chiao, Bemjamin, Lerner, Josh, & Tirole, Jean. 2007. The Rules of Standard Setting Organizations: An Empirical Analysis. *RAND Journal of Economics*, **38**(4), 905–930.

David, Paul A., & Greenstein, Shane. 1990. The Economics of Compatibility Standards: An Introduction to Recent Research. *Economics of Innovation and New Technology*, **1**, 3–42.

Davies, Elwyn. 2004 (May). *IETF Problem Statement*. Request For Comments 3774. Internet Engineering Task Force.

Farrell, Joseph, & Saloner, Garth. 1988. Coordination through Committees and Markets. *RAND Journal of Economics*, **19**(2), 235–252.

Farrell, Joseph, & Shapiro, Carl. 1992. Standard Setting in High-Definition Television. *Brookings Papers on Economic Activity*, 1–93.

Farrell, Joseph, & Simcoe, Timothy. 2008. Choosing the Rules for Formal Standardization. *Unpublished manuscript*.

Farrell, Joseph, Hayes, John, Shapiro, Carl, & Sullivan, Theresa. 2007. Standard Setting, Patents and Hold-Up. *Antitrust Law Journal*, **74**, 603–670.

Foray, Dominique. 1995. *Standards, Innovation and Competitiveness*. Edward Elgar. Chap. Coalitions and Committees: How users get involved in information technology standardization, pages 192–212.

Furman, Jeff, & Stern, Scott. 2006. Climbing Atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research. *NBER Working Paper No. 12523*.

Hall, Bronwyn, Jaffe, Adam, & Trajtenberg, Manuel. 2005. Market Value and Patent Citations. *RAND Journal of Economics*, **36**, 16–38.

Heckman, James. 1979. Sample selection bias as a specification error. *Econometrica*, **47**, 153–161.

IETF. 1990 (June). *Internet Monthly Report*. Tech. rept. Internet Engineering Task Force.

Kennan, John, & Wilson, Robert. 1993. Bargaining with Private Information. *Journal of Economic Literature*, **31**(1), 45–104.

Lee, Lung-Fei. 1978. Unionism and Wage Rates: A Simultaneous Equations Model with Qualitative and Limited Dependent Variables. *International Economic Review*, **19**(2), 415–433.

Lehr, William. 1995. Compatibility Standards and Interoperability: Lessons from the Internet. *In:* Kahin, Brian, & Abbate, Janet (eds), *Standards Policy for Information Infrastructure*. Cambridge, MA: MIT Press.

Leiponen, Aija. 2008. Competing Through Cooperation: Standard Setting in Wireless Telecommunications. *Management Science*, **54**(2), 1904–1919.

Lemley, Mark. 2002. Intellectual Property Rights and Standard Setting Organizations. *California Law Review*, **90**, 1889–1981.

Lemley, Mark. 2007. Ten Things To Do About Patent Holdup of Standards (And one Not To). *Boston College Law Review*, **48**, 149–168.

Lerner, Josh, & Tirole, Jean. 2006. A Model of Forum Shopping. *American Economic Review*, **96**(4), 1091–1113.

Lippman, Steven A., & McCall, John J. 1976. The Economics of Job Search: A Survey. *Economic Inquiry*, **4**, 17–37.

Maddala, G.S. 1983. *Limited-Dependent and Qualitative Variables in Econometrics*. Cambridge University Press.

Merlo, Antonio, & Wilson, Charles. 1995. A Stochastic-Model of Sequential Bargaining with Complete Information. *Econometrica*, **63**(2), 371–399.

Myatt, David P. 2005. Instant Exit from the Asymmetric War of Attrition. *Oxford University Economics Working Papers*, **No. 160**.

National Research Council. 1990. *Crossroads of Information Technology Standards*. National Academy Press.

Postel, Jon. 1995. *Internet Official Protocol Standards (RFC 1800)*. Tech. rept. Internet Engineering Task Force.

Rosenbaum, Paul R., & Rubin, Donald B. 1983. The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika*, **70**, 41–55.

Rosenkopf, Lori, Metiu, Anca, & George, Varghese P. 2001. From the Bottom Up? Technical Committee Activity and Alliance Formation. *Administrative Science Quarterly*, **46**, 748–772.

Rysman, Marc, & Simcoe, Timothy. 2008. Patents and the Performance of Voluntary Standard Setting Organizations. *Management Science*, **54**(11), 1920–1934.

Saltzer, Jerome H., Reed, David P., & Clark, David D. 1984. End-to-End Arguments in System Design. *ACM Transactions on Computer Systems*, **2**(4), 277–288.

Shapiro, Carl, & Varian, Hal R. 1998. *Information Rules : A Strategic Guide to the Network Economy*. Boston, Mass.: Harvard Business School Press.

Updegrove, Andrew. 2007. Intellectual Property Rights and Standard Setting. *Consortium Stadnards Bulletin*, **6**(3).

Waguespack, David, & Fleming, Lee. 2008. Scanning the commons? Evidence on the benefits to statrups of participation in open standards development. *Management Science*, **forthcoming**.

Weiss, Martin B., & Sirbu, Marvin. 1990. Technological Choice in Voluntary Standards Committees: An Empirical Analysis. *Economics of Innovation and New Technology*, **1**, 111–134.

Yildiz, Muhamet. 2004. Waiting to Persuade. *The Quarterly Journal of Economics*, **119**(1), 223–248.

Table 1: Sample Size & Average Duration by Draft Cohort

	IETF Activity		Observations (Count) by Publication-type				Mean Duration (Total Days)		
Draft Cohort	Total Drafts	Active WGs†	S-Track	N-Track	Expired	Censored	S-Track	N-Track	Expired
1993	58	35	25	19	14	0	490	221	174
1994	86	41	43	19	24	0	514	484	419
1995	128	48	67	21	40	0	738	426	353
1996	167	61	87	39	41	0	751	535	660
1997	304	76	119	57	128	0	673	558	472
1998	246	73	78	55	113	0	488	652	447
1999	279	82	94	71	113	1	814	652	503
2000	326	79	143	55	125	3	930	757	469
2001	379	100	131	76	159	13	1002	692	598
2002	325	87	139	70	98	18	847	592	423
2003	303	85	127	71	78	27	703	514	491
Total	2,601	249	1,053	553	933	62	774	595	487

†Active = One or more new Internet Drafts submitted.

Table 2: Variable Definitions & Summary Stats

Variable Name	Definition	Variation†	Mean	S.D.
	Outcome Measures			
Total Days	Days from initial to final ID submission	ID	659.75	566.97
Versions	Count of ID revisions submitted	ID	6.33	4.39
S-Track	Dummy for Proposed Standard	ID	0.40	0.49
N-Track	Dummy for Nonstandards-track RFC	ID	0.21	0.41
log(ID Mail)	Log count of emails mentioning ID	ID	2.98	1.25
RFC Cites	Cites from other RFCs	RFC	7.93	15.25
Patent Cites	US Patent non-patent prior art cites	RFC	2.28	9.54
Article Cites	ISI academic journal cites	RFC	1.65	7.17
	Distributional Conflict (ω)			
Suit-share	% of all domains on WG email listserv in past year from dot-com or dot-net TLDs‡	WG	73.34	16.04
log(Patents)	Log 5-year patent stock of all ID authors	WG	7.61	2.98
	Access to Concessions (γ)			
Dot-org	Dummy for author from dot-org TLD	ID	0.08	0.26
Dot-edu	Dummy for author from dot-edu TLD	ID	0.18	0.39
Dot-gov	Dummy for author from dot-gov TLD	ID	0.05	0.21
	Control variables			
Draft Cohort	Initial submission year	ID	1999.24	2.68
log(WG Mail)	Log past-year e-mail messages	WG	5.86	1.71
log(Cum Mail)	Log total e-mail messages	WG	6.86	1.83
log(Drafts)	Log drafts under review	WG	1.89	0.71
log(Cum Drafts)	Log total drafts submitted	WG	2.55	1.06
log(Members)	Log orgs (TLDs) with submissions	WG	2.19	0.84
log(Filesize)	Log size (Bytes) of initial draft	ID	10.29	0.89
log(Keywords)	Log count of keywords	ID	1.74	1.54
ID Sponsors	Count of author affiliations	ID	1.97	1.32
WG Chair	Dummy for past WG chair author	ID	0.42	0.49
Dot-com	Dummy for author from dot-com TLD	ID	0.86	0.35
ID Suit-share	% of all emails that mention focal ID from dot-com or dot-net TLDs	ID	67.69	22.28
Applications Area	IETF Technology Class dummy	WG	0.18	0.38
Transport Area	IETF Technology Class dummy	WG	0.20	0.40
Internet Area	IETF Technology Class dummy	WG	0.23	0.42
Routing Area	IETF Technology Class dummy	WG	0.12	0.32
Operations Area	IETF Technology Class dummy	WG	0.14	0.34
Security Area	IETF Technology Class dummy	WG	0.13	0.34

† WG = Working Group; ID = Internet Draft; RFC = Request for Comments. ‡ TLD = Top-level domain (e.g. Dot-com or Dot-edu).

Table 3: Sample Means for Standards and Nonstandards

	All RFCs			Matched Sample†		
Variable	Nonstds Track	Proposed Standards	P-value (ST=NST)	Nonstds Track	Proposed Standards	P-value (ST=NST)
Total Days	607.50	784.01	0.00	544.82	703.99	0.00
Versions	5.99	8.19	0.00	5.45	7.63	0.00
log(Keywords)	1.36	1.88	0.00	1.42	1.83	0.00
log(ID Mail)	3.03	3.43	0.00	3.17	3.50	0.00
RFC Cites	4.89	10.41	0.00	5.31	9.84	0.00
Patent Cites	1.72	3.04	0.01	1.79	2.99	0.07
Article Cites	1.58	1.97	0.41	1.62	1.96	0.63
Suit-share	71.62	73.62	0.05	73.38	73.29	0.94
log(Patents)	7.39	7.26	0.47	7.12	7.45	0.15
Dot-org	0.08	0.09	0.67	0.08	0.07	0.53
Dot-edu	0.22	0.16	0.01	0.18	0.17	0.71
Dot-gov	0.06	0.05	0.43	0.06	0.06	0.98
Draft Cohort	1998.66	1998.58	0.56	1998.65	1998.70	0.77
log(WG Mail)	5.45	5.86	0.00	5.59	5.84	0.05
log(Cum Mail)	6.29	6.85	0.00	6.49	6.85	0.00
log(Drafts)	1.68	1.85	0.00	1.82	1.93	0.01
log(Cum Drafts)	2.21	2.51	0.00	2.44	2.64	0.00
log(Members)	2.13	2.16	0.56	2.11	2.19	0.11
log(Filesize)	10.24	10.34	0.06	10.19	10.26	0.23
ID Sponsors	2.14	1.99	0.06	1.99	2.00	0.92
WG Chair	0.43	0.46	0.33	0.41	0.44	0.37
Dot-com	0.85	0.88	0.21	0.86	0.87	0.55
Hotelling's T^2 (P-value)	0.00			0.12		
Obs. (Standards)	926			671		
Obs. (Nonstds)	482			287		

Notes: Left panel includes all RFCs with Draft-Cohort $\leq$ 2002. Right panel omits RFCs with Total Days > 2007 (5.5 years) to correct for right-truncation in OLS models. †See appendix for details on propensity-score matching.

Table 4: Distributional Conflict and Publication Delays

	Dependent Variable = Total Days				
Model	OLS Diffs (1)	Matched Diffs (2)	WG FE Diffs (3)	OLS Standards (4)	WG FE Standards (5)
Suit-share * S-track	8.3 (2.2)***	8.1 (2.4)***	7.2 (2.5)***	4.8 (1.5)***	4.7 (2.4)*
log(Patents) * S-track	12.0 (9.8)	6.7 (10.5)	5.0 (11.7)	15.4 (7.2)**	18.8 (9.6)*
Dot-gov * S-Track	76.5 (67.9)	44.7 (79.0)	142.7 (93.8)	12.2 (65.6)	59.5 (75.1)
Dot-org * S-Track	223.8 (79.9)***	212.2 (94.2)**	197.1 (105.8)*	142.1 (65.7)**	173.0 (80.1)**
Dot-edu * S-Track	164.4 (62.4)***	139.4 (78.1)*	140.5 (77.8)*	81.7 (39.3)**	63.3 (50.2)
Suit-share	-3.7 (1.8)**	-2.0 (1.9)	-0.9 (2.8)		
log(Patents)	2.1 (7.6)	0.3 (8.1)	7.4 (11.1)		
Dot-gov/org/edu	-100.4 (54.1)*	-55.0 (72.8)	-84.5 (64.1)		
	Controls & Regression Statistics				
Additional Controls†	8 [0.00]	8 [0.00]	8 [0.00]	4 [0.00]	4 [0.00]
Draft Cohort FEs	18 [0.00]	18 [0.00]	18 [0.01]	9 [0.00]	9 [0.00]
Technology Class FEs	5 [0.00]	5 [0.01]		5 [0.00]	
Working Group FEs			152 [0.00]		132 [0.00]
Hausman $\chi^{2\dagger}$			9.29 [0.99]		29.73** [0.04]
Obs. (Standards)	772	671	671	772	772
Obs. (Nonstandards)	392	287	287	0	0
Model dof	40	40	186	23	149
R-Squared	0.19	0.20	0.11	0.21	0.09

*10% significance; **5% significance; ***1% significance (all SEs clustered on Working Group). **Notes**: All models omit observations with Total Days > 2007 (5.5 years) to correct for right-truncation of the DV. The sample in (1), (4) and (5) contains all standards (nonstandards) from Draft-Cohorts 1993 through 2002. Models (2) and (3) use a propensity-score matched sample (see online appendix for details). For controls, the table indicates number of parameters and a p-value for joint-significance. †Additional controls are main effects and interactions for: log(Drafts), log(WG Mail), ID Sponsors, log(Filesize). See the text for additional details.

Table 5: Conflict and Delays: Switching Model

Model	Switching Regressions (1)			OLS (2)	OLS (3)
Dependent Variable	S-Track log(Days)	N-Track log(Days)	S-Track Selection	S-Track log(Days)	N-Track log(Days)
log(Suit-share)	0.362	-0.534	0.220	0.378	-0.421
	(0.169)**	(0.184)***	(0.262)	(0.146)**	(0.186)**
log(Patents)	0.009	0.020	-0.027	0.025	-0.006
	(0.017)	(0.022)	(0.021)	(0.014)*	(0.018)
Dot-org	0.037	-0.211	0.075	0.201	0.045
	(0.108)	(0.190)	(0.16)	(0.107)*	(0.164)
Dot-edu	0.051	0.061	-0.182	0.092	-0.115
	(0.064)	(0.100)	(0.114)	(0.076)	(0.101)
Dot-gov	-0.180	-0.202	-0.242	-0.138	-0.260
	(0.157)	(0.237)	(0.195)	(0.163)	(0.155)*
log(Keywords)			0.140	-0.022	0.042
			(0.034)***	(0.024)	(0.029)
Error corr. $(\sigma_{\nu x}/\sigma^2_{\varepsilon^x})$	0.258	-0.067			
	(0.087)***	(0.210)			
	Controls & Regression Statistics				
Additional Controls†	12 [0.00]			4 [0.00]	4 [0.02]
Technology Class FEs	10 [0.00]			5 [0.01]	5 [0.18]
Draft Cohort FEs	30 [0.00]			9 [0.00]	9 [0.00]
Obs. (Standards)	935			772	0
Obs. (Nonstandards)	474			0	392
Obs. (Expired)	820				
Obs. (Censored)	58				
	Joint Hypothesis Tests				
Suit-to-Beard Equality	$\chi^2(1) = 11.74$***				
log(IPR) Equality	$\chi^2(1) = 0.15$				
Endogeneity $(\rho \neq 0)$	$\chi^2(2) = 10.03$***				

*10% significance; **5% significance; ***1% significance (all SEs clustered on Working Group). **Notes**: The sample for model (1) contains all IDs submitted between 1993 and 2003. OLS models (2) and (3) omit observations with Total Days > 2007 (5.5 years) to correct for right-truncation of the DV. For controls, the table indicates number of parameters and a p-value for joint significance. ‡Additional controls are log(Drafts), log(WG Mail), Sponsors, and log(Filesize). See the text for a discussion of the switching model.

Table 6: Conflict and Delay: Time and Technology Interactions

	Dependent Variable = Total Days			
Model	OLS (1)	OLS (2)	OLS (3)	OLS (4)
Applications Area†	5.5 (2.5)**	5.8 (2.5)**		
Transport Area†	6.4 (2.5)**	6.5 (2.5)**		
Internet Area†	3.4 (2.3)	3.8 (2.3)*		
Routing Area†	9.7 (2.3)***	9.9 (2.4)***		
Security Area†	5.8 (2.6)**	5.9 (2.6)**		
Operations Area†	8.9 (2.3)***	9.1 (2.3)***		
Cohort [93,94]†			2.9 (3.9)	-1.2 (4.7)
Cohort [95,96]†			6.6 (2.9)**	5.2 (4.2)
Cohort [97,98]†			5.9 (3.3)*	7.3 (4.2)*
Cohort [99,00]†			10.8 (2.9)***	7.8 (3.6)**
Cohort [00,01]†			9.8 (3.5)***	12.6 (3.6)***
	Controls & Regression Statistics			
Suit-share * Technology FEs		6 [0.41]		
Suit-share * Cohort FEs				5 [0.27]
Technology Class FEs			Y	Y
Draft Cohort FEs	Y	Y	Y	Y
Additional Controls‡	Y	Y	Y	Y
Working Group FEs	Y	Y	N	N
Obs. (Standards)	772	772	772	772
Obs. (Nonstandards)	392	392	392	392
Model d.o.f.	213	218	44	48
R-Squared	0.11	0.11	0.19	0.19

*10% significance; **5% significance; ***1% significance (all SEs clustered on Working Group). **Notes**:†Indicates coefficient for a the three-way interaction between listed variable, S-track dummy and Suit-share. The sample in (1) through (4) contains all standards and nonstandards from Draft-Cohorts 1993 through 2002 with Total Days $\leq$ 2007. ‡Additional controls are main effects and standards-track interactions for: log(Drafts), log(WG Mail), Sponsors, log(Filesize), Dot-edu, Dot-gov and Dot-org. See the text for additional details.

Table 7: Publication Delay and Forward Citations

Model	Poisson (1)	Poisson (2)	Poisson (3)	Poisson (4)	Poisson (5)	Poisson (6)
Dependent Variable	All Cites	All Cites	All Cites	RFC Cites	Patent Cites	Article Cites
log(Suit-share)	0.22 (0.34)			0.65 (0.40)	1.42 (0.90)	-1.16 (0.26)***
log(Suit-share) * S-track	-0.18 (0.33)			-0.72 (0.40)*	-1.00 (0.88)	1.16 (0.22)***
log(Total Days)		-0.05 (0.08)				
log(Total Days) * S-track		0.32 (0.11)***				
log(Versions)			-0.11 (0.13)	-0.11 (0.17)	0.23 (0.28)	-0.20 (0.20)
log(Versions) * S-track			0.61 (0.13)***	0.48 (0.18)***	0.60 (0.30)**	0.49 (0.20)**
log(Cites/Month)$_{-i}$	0.63 (0.05)***	0.62 (0.05)***	0.60 (0.06)***	0.54 (0.06)***	0.81 (0.20)***	0.76 (0.12)***
log(ID Mail)	0.33 (0.03)***	0.33 (0.03)***	0.27 (0.02)***	0.29 (0.03)***	0.14 (0.06)**	0.35 (0.07)***
	Control Variables & Regression Statistics					
Draft Cohort FE's	10 [0.00]	10 [0.00]	10 [0.00]	10 [0.00]	10 [0.00]	10 [0.00]
RFC Month Polynomial	4 [0.00]	4 [0.00]	4 [0.00]	4 [0.00]	4 [0.00]	4 [0.00]
Additional Controls	12 [0.00]	12 [0.00]	12 [0.00]	12 [0.00]	12 [0.00]	12 [0.00]
S-track x Tech-Class FEs	Y	Y	Y	Y	Y	Y
Obs. (Standards)	860	860	860	860	860	860
Obs. (Nonstandards)	435	435	435	435	435	435
Model d.o.f.	29	29	29	31	31	31
Pseudo-LogL x 10^{-3}	-8.45	-8.35	-8.23	-6.39	-2.53	-2.28

*10% significance; **5% significance; ***1% significance (all SEs clustered on Tech Class by S-Track).
Notes: The sample in models (1) through (6) is all RFCs from Draft Cohorts 1993 through 2003 published before 2007. For controls, the table indicates number of parameters and a p-value for joint significance. Additional controls are main effects and standards-track interactions for: log(Drafts), log(WG Mail), Sponsors, log(Filesize), Dot-edu, Dot-gov and Dot-org. See the text for additional details.

Appendix A: Proofs

The text contains simple proofs of P2 and P3. The following two facts are useful for proving P1, P4 and P5. First, substituting $p \equiv 1[x \geq q]$ into the Bellman equation (1) provides an expression for $V(q)$:

$$V = -c + \frac{\int_q^\infty \left\{E\left[\pi(x,a)\right] + \frac{k-1}{k}(\gamma-1)b(x)\right\} dG(x)}{1-\beta G(q)} \tag{A-1}$$

And second, since (2) and (3) both bind at the equilibrium stopping threshold, we can set them equal to solve for the maximum concessions:

$$b(q^*) = \frac{\pi(q^*,1) - \pi(q^*,0)}{k-1+\gamma} \tag{A-2}$$

Proof if P1: *There is a unique symmetric equilibrium.*
Proof: Given V, there is a unique q^* that solves equation (4). To see this, note that the right side of (4) is strictly decreasing and therefore crosses kc from above exactly once between $\underline{q}$ and ∞. Thus, there cannot be two symmetric equilibria with the same V.

Now suppose there is a pair of symmetric equilibria where $V' > V$. This implies more concessions in the high-value equilibrium: $b'(x) \geq b(x)$ for all $x \geq q^*$. And since concessions increase the last term in (4), the high-value equilibrium also has a higher stopping threshold $q' > q^*$. Differentiation shows that (A-1) is decreasing in q for all $q > q^*$. So increasing both q^* and b leads to $V' < V$, a contradiction. $\square$

Proof of P4: *Inefficient delays are decreasing in γ.*
Proof: First, note that γ has no impact on the optimal stopping rule q^{fb}. Thus, a decrease in q^* implies a reduction in $G(q^*) - G(q^{fb})$, which is equivalent to a decline in (expected) inefficient delays. Totally differentiating (4) with respect to γ yields:

$$\frac{dq^*}{d\gamma} = \frac{(k-1)\left\{\int_{q^*}^\infty b(x)dG(x) - \Delta(q^*)b(q^*)\right\} + (k-1)(\gamma-1)\left\{\int_{q^*}^\infty b_\gamma(x)dG(x) - \Delta(q^*)b_\gamma(q^*)\right\}}{\left[S_q(q^*) + (k-1)(\gamma-1)b_q(q^*)\right]\Delta(q^*)}$$

Since $S(x)$ is increasing and $b(x)$ decreasing, the denominator is positive and the first term in the numerator is negative. Thus, it suffices to show that the second term in the numerator is negative.

From (A-2), we have $b_\gamma(q^*) = \frac{-b(q^*)}{k+1-\gamma}$. Differentiating $b(x)$ reveals that $b_\gamma(x) = \frac{\beta V_\gamma - b(x)}{\gamma}$ (for all $b(x) > 0$). Plugging in these expressions (and recalling that $\Delta(q^*) \equiv \beta^{-1} - G(q^*)$) shows that the curly-bracketed part of the final term in the numerator equals

$$\frac{1}{\gamma}\left(\int_{q^*}^{\infty}[b(q^*)-b(x)]\,dG(x)+\frac{1-\beta}{\beta}\frac{b(q^*)}{k+\gamma-1}\right)>0$$

so q^* is decreasing in γ. $\square$

Proof of P5: *As k grows large, q^* approaches q^{fb}.*

Proof: Divide both sides of (4) by k and consider what happens as $k \to \infty$. Since there are $k-1$ losers and one winner the average gross-payoff converges towards that of a loser: $\frac{S(x)}{k} \to \pi(x,0)$. Substituting (A-2) into the final term of (4) shows that:

$$0<\frac{(\gamma-1)(k-1)}{k}\left\{\int_{q^*}^{\infty}b(x)dG(x)-\Delta(q^*)b(q^*)\right\}<\frac{(1-\gamma)(k-1)\Delta(q^*)[\pi(q^*,1)-\pi(q^*,0)]}{k(k-1+\gamma)}$$

Since the last term in this inequality approaches zero as k grows large, average concessions (the middle term) must also converge to zero. This implies that $q^* \to q^{fb}$; the optimal stopping rule for a single agent who is certain to receive a losing payoff. $\square$

Appendix B: Data Set Construction

Dependent Variables

Time-to-Consensus

Data on every Internet Draft published between 1993 and 2003 was obtained from two sources: 1) the "ietf-announce" mailing list, and 2) www.watersprings.org. Publication dates come from the ietf-announce list in over 90 percent of the cases, and the watersprings site was used to fill in dates when one version in a particular series was missing.

Citations

I collected citation data from three sources. RFC citations were gather by using a Perl program to examine the reference section of an ASCII text copy of every published RFC. The complete RFC archives are available at www.rfc-editor.org. Patent citations were collected from the USPTO website, where I searched for all non-patent prior art citations containing the string "RFC" or "Request for Comments" followed by a four digit number. Finally, academic journal citations were collected by performing a similar cited-reference search using all journals in the ISI Web of Science database.

Independent Variables

Internet Drafts

As explained in the text, the IETF's file naming conventions were used to match each Internet Draft to a particular Working Group. Author attributes were collected by using a Perl script to parse the header and acknowledgements section of each ID. Similarly, key word counts were obtained by using Perl to scan an ASCII text copy of the proposal.

E-mail Addresses and Working Group Discussion Lists

Data on committee demographics were obtained from the archived e-mail discussion lists of 278 IETF Working Groups. Many of these can be found at ftp.ietf.org/ietf-mail-archive, and the remainder were located by searching the Internet. Collectively, these data go back as far as the late 1980's, and comprise more than 690,000 messages. Most Working Group communication takes place on these e-mail discussion lists.

I used Perl scripts to examine each message and construct measures of WG demographics, participation, and experience. All of these variables are based on information contained in message headers—specifically the date, sender's address, and subject line fields. I used a number of well known domain-naming conventions to classify the institutional affiliation of users from less common top-level domains (e.g. bt.co.uk, rotman.utoronto.ca, or alvestrand.no).

I used several approaches to address the problems of spam and hosted mail. First, I exclude all messages with subject lines related to pornography, home mortgages, hot stocks, or exciting new business opportunities$!$!$! Second, I removed messages originating from the most popular hosted mail sites (e.g. Yahoo! and Hotmail). While this criteria may drop some legitimate messages, I found that most IETF participants have several e-mail accounts—one of which was generally within the domain of their employer. Finally, the results in the paper are based on a sample of messages whose sender (originating address) appeared four or more times on the same list on different dates with different subject lines. I also constructed a sample based on messages that were part of a discussion thread, i.e. either generated a reply or replied to an earlier message. All of the results are robust to a variety of changes in these rules and the criteria used to screen messages.

The e-mail data can be aggregated at three levels: message, sender (unique address), or organization (unique top-level domain). For all variables, the correlation across these different levels is extremely high, e.g. Suit-share measures based on messages, participants, and firms are all correlated above 0.98. Consequently, the results do not change if I change the aggregation level for a particular variable, but they do become unstable when I try to include all three levels.

The NBER Patent Data Merge

The "affiliation" of each Internet Draft author is identified using their email address. This approach identifies 1,460 organizations with one or more contributions to an IETF Working

Group. I focus on 498 organizations that appear on two or more Internet Drafts, and attempted to match each organization to a standardized patent assignee number by hand.[35]

I successfully matched 193 organizations to an assignee code. All of the top 100 IETF contributors were either matched or determined to be non-US patent holders. Many of the unmatched organizations were non-profits (e.g. the World Wide Web Consortium), network operators (e.g. MERIT) or non-US academic institutions. Because of the skew in contribution rates, one or more of the matched organizations appeared as an author on over 90 percent of the proposals in the estimation sample. However, some small patenting firms may not show up in the data because of lags in the US Patent and Trademark Office's patent review and reporting process.

For each firm in the matched group, I construct a five year unadjusted patent stock, and a five year stock depreciated at 15 percent. I also calculated the uncentered correlation coefficient over 3-digit USPTO technology classes between the firm's patent portfolio and the total stock of patents owned by IETF participants. The log(Patents) variable used in the analysis is the sum (over all firms with one or more proposals before a Working Group) of the five-year depreciated patent stock weighted by the firm-specific correlation coefficient.

[35]See http://elsa.berkeley.edu/ bhhall/pat/namematch/namematch.html for a discussion of assignee names and numbers.

NOT FOR PUBLICATION APPENDIX: Additional Material for Referees

Propensity Score Matching

I use fitted values $\widehat{P(x)}$ from a probit model of standards-track assignment to estimate the propensity score. Estimates from this probit model are presented in Table R-1. Figure R-1 shows the empirical distribution of the estimated propensity-score for standards and nonstandards in the estimation sample. The solid vertical lines correspond to the 5th percentile of the estimated propensity score distribution for Proposed Standards (P_5^s) and the 95th percentile of the propensity score distribution for nonstandards-track RFCs (P_{95}^n). (The dotted lines correspond to the 1st and 99th percentiles.) The region where $P_5^s \leq \widehat{P(X_i)} \leq P_{95}^n$ is the common support of the p-score distribution. Discarding observations that fall outside this range (i.e. "trimming" or "blocking" on the propensity score) leaves 958 out of the original 1,164 RFCs, or 82 percent of the initial sample.

Figure R-1: Propensity Score Distribution

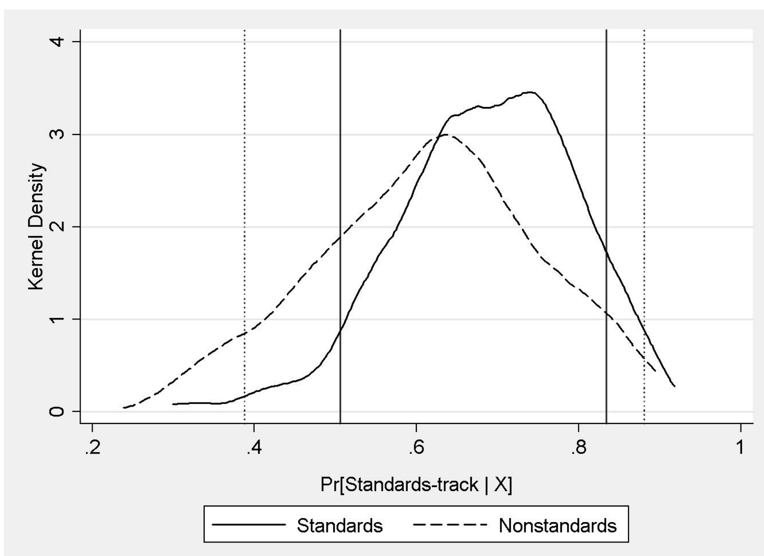

Table R-1: Propensity Score Probit

	Probit Regression Marginal Effects and SEs
	DV = 1[Standards-Track]
Suit-share	-0.00
	(0.00)**
ID Suit-share	0.00
	(0.00)***
log(Patents)	-0.02
	(0.01)**
log(Drafts)	0.02
	(0.05)
log(Cum Drafts)	0.03
	(0.04)
log(Members)	0.04
	(0.03)
log(Email)	-0.00
	(0.02)
log(Cum Email)	0.02
	(0.02)
Sponsors	-0.02
	(0.01)*
WG Chair	0.07
	(0.03)**
log(Filesize)	0.05
	(0.02)***
Dot-org	0.07
	(0.05)
Dot-edu	-0.06
	(0.04)
Dot-gov	-0.03
	(0.06)
Tech Class Effects	5
ID Cohort Effects	2
WG Cohort Effects	2
Observations	1164
Model dof	27

Model includes technology class effects and quadratic in draft-year and WG cohort.

Switching Model Log-Likelihood

The switching model is defined by Equations (5), (6) and (8). The correlation between ε^s and ε^n is undefined, since we never observe both T^s and T^n. Following the approach described in Lokshin and Sajaia (2004), I define

$$\eta_{is} = \frac{\gamma Z_i + \rho_{sv}\varepsilon_i^s/\sigma_s}{\sqrt{1-\rho_{sv}^2}}$$

where $\rho_{sv} = \sigma_{sv}/\sigma_s$ is the coefficient of correlation between ε^s and ν. Since the ε have a trivariate normal distribution, there is a simple closed form expression for uncensored observations (i.e. Proposed Standards and nonstandards-track RFCs).

$$Pr(S_i = 1, T_i^s = T) = \frac{\Phi(\eta_{is})}{\sigma_s}\phi\left(\frac{\varepsilon_i^s}{\sigma_s}\right)$$
$$Pr(S_i = 0, T_i^n = T) = \frac{(1-\Phi(\eta_{in}))}{\sigma_n}\phi\left(\frac{\varepsilon_i^n}{\sigma_n}\right)$$

I treat expired proposals as censored observations whose intended publication status is unknown. Since no proposal is on both the standards- and nonstandards-track, the probability of observing a censored or expired proposal must be

$$Pr(S_i = 1, T_i^s > T_i) + Pr(S_i = 0, T_i^n > T_i) = \Phi\left(\gamma Z_i, \frac{-\varepsilon_i^s}{\sigma_s}, \rho_{sv}\right) + \Phi\left(-\gamma Z_i, \frac{-\varepsilon_i^n}{\sigma_n}, -\rho_{nv}\right)$$

where Φ is the cumulative bivariate normal distribution with correlation parameter ρ.

The log-likelihood is just $\sum_i \ln(Pr(S_i, T_i))$. Code for estimating this model in Stata was adapted from the movestay routine developed by Lokshin and Sajaia (2004) and is available from the author.

Additional Descriptive Statistics

Table R-2: Top IETF Contributors† (1992-2004)

1992-1994		1992-2004	
1. Cisco	94	1. Cisco	1,787
2. Carnegie Mellon	51	2. Nortel	694
3. mtview.ca.us	48	3. Microsoft	581
4. IBM	44	4. Nokia	539
5. SNMP Research	38	5. Sun Microsystems	513
1995-1997			
1. Cisco	214	6. AT&T	513
2. IBM	140	7. IBM	490
3. Microsoft	140	8. Ericsson	398
4. Sun Microsystems	84	9. Lucent	343
5. USC (ISI)	79	10. Bell Labs	301
1998-2000			
1. Cisco	517	11. Alcatel	299
2. Nortel	321	12. Juniper Networks	260
3. AT&T	223	13. Intel	225
4. Microsoft	221	14. Columbia U.	220
5. Sun Microsystems	180	15. Siemens	200
2001-2004			
1. Cisco	962	16. Dynamicsoft	196
2. Nokia	404	17. USC (ISI)	195
3. Nortel	354	18. ACM	185
4. Ericsson	279	19. MIT	152
5. Sun Microsystems	234	20. NTT	149

†Rankings are based on the number of Internet Drafts submitted during the relevant period.

Table R-3: Sample Construction

	Individual Drafts	Working Group Drafts†	One or more Revisions	Estimation Sample‡
	Total Observations			
Total Internet Drafts	6,481	3,521	2,662	2,601
Working Groups	0	283	272	249
	Internet Draft Outcomes			
Censored	47	62	62	62
Expired	5,729	1,730	950	933
Nonstandards-track RFC	534	641	586	553
Proposed Standard	171	1,088	1,064	1,053

†Excludes IDs from General and User Service Areas, IDs originating before 1993 or after 2003, BCP's, Historic RFC's, Draft Standards and Internet Standards.

Table R-4: WG Publication (RFC) Counts†

Proposed Standards	Nonstandards-track RFCs 0	1	2	3	4	≥5	Total
0	25	27	9	4	2	6	73
1	21	12	5	3	1	4	46
2	9	11	1	2	0	4	27
3	8	6	1	2	1	3	21
4	7	4	2	2	0	0	15
≥5	9	11	15	4	5	23	67
Total	79	71	33	17	9	40	249

† Each cell contains a count of WGs that published the number of standards (nonstandards) indicated by the row (column) headings.

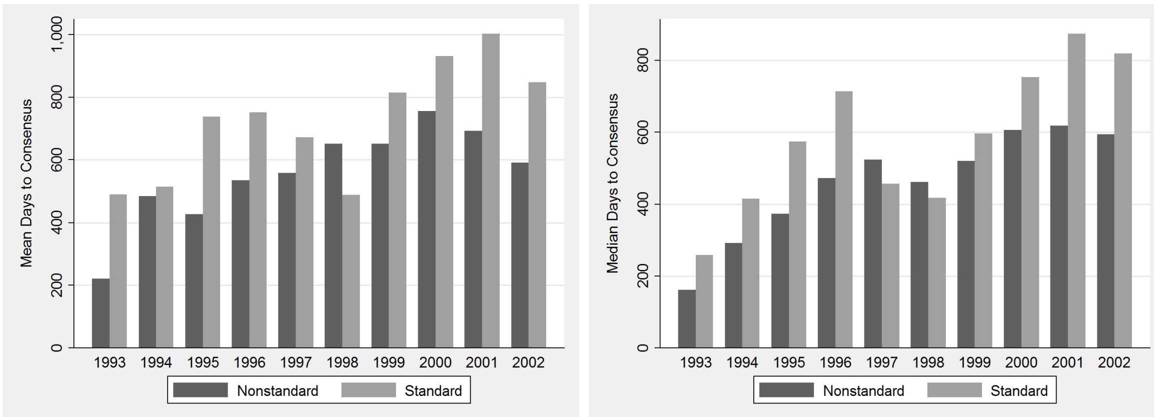

Figure R-2: RFC Duration by Draft-Cohort: Mean (left) & Median(right)

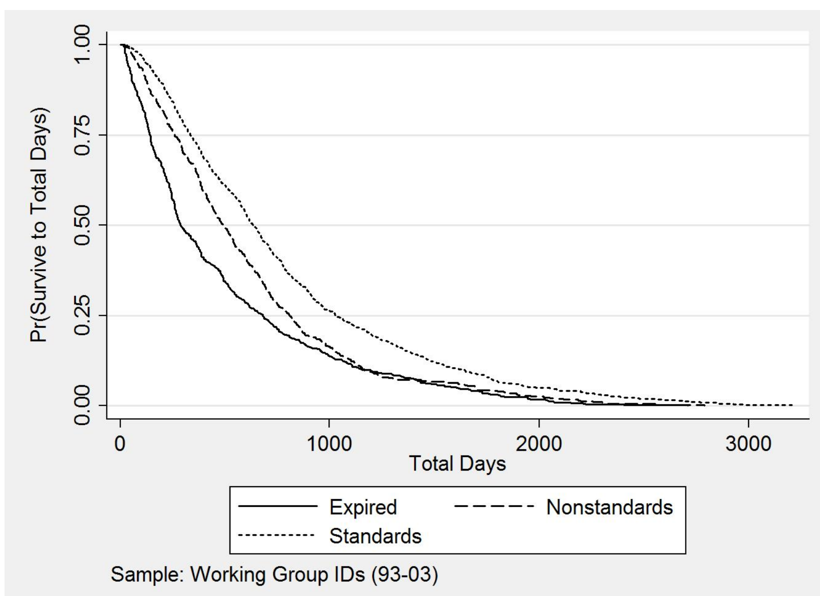

Figure R-3: Kaplan-Meier Survival Curves by Internet Draft Outcome

Robustness Checks

Table R-5: Instrumenting for Suit-share

	Obs = Internet Draft DV = Total Days			Obs = WG-Year-Track DV = Avg. Total Days			
Model	GMM IV (1)	GMM IV (2)	GMM IV (3)	OLS (4)	OLS (5)	Arellano Bond (6)	Arellano Bond (7)
Suit-share * S-track	6.9	10.8	14.9	4.4	3.4	8.5	8.6
	(2.3)***	(2.9)***	(12.0)	(2.0)**	(4.8)	(6.4)	(9.7)
Suit-share		-4.8		-1.9	-0.4		-0.5
		(2.6)*		(2.8)	(4.2)		(7.9)
	First-Stage Statistics						
Instruments							
Lag-Suit-share	Y	Y	N				
Lag-Share * S-track	N	Y	N				
Non-WG Tech Area Growth	N	N	Y				
Partial R-squared							
Suit-share * S-track	0.43 [0.00]	0.47 [0.00]	0.04 [0.06]				
Suit-share		0.40 [0.00]					
	Controls & Regression Statistics						
Tech Class FEs	Y	Y					
PubCohort FEs	Y	Y	Y	Y	Y	Y	Y
Working Group FEs	N	N	Y	Y			
WG-Track FEs	N	N	N	N	Y	Y	Y
Additional Controls	Y	Y	Y	Y	Y	N	N
Obs (S-track)	671	671	737	334	334	91	91
Obs (N-track)	0	287	0	245	245	0	43
Model F / χ^2	13.23***	10.83***	7.12***	2.12**	0.79	16.15*	13.90

*10% significance; **5% significance; ***1% significance (all SEs clustered on Working Group). **Notes**: Models (1) through (3) take the RFC as a unit of observation and instrument for Suit-share using lag Suit-share or the growth rate of other WGs in the same Technology Area. Models (4) through (7) collapse the data to WG-Year-Track level. All models omit observations with Total Days > 2007 (5.5 years) to correct for right-truncation of the DV. Models(1) and (2) use the matched sample for comparison to Table 4. Additional controls in (1) through (3) are main effects and S-track interactions for log(Filesize), Sponsors, log(WG Mail), log(Drafts) and Dot-com/org/edu. Additional controls in (4) and (5) are log(Drafts), log(WG Mail), log(Cum Drafts), log(Cum Mail) and log(Members).

Table R-6: Diff-in-diffs: Alternative Measures and Specifications

Model	Big-Firm Effects		ID Suit-Share		DV = Versions		Cox Hazard	
Suit-share * S-track	8.40	7.33			0.05	0.05	-1.56	-1.96
	(2.43)***	(2.56)***			(0.02)**	(0.02)**	(0.49)***	(0.51)***
log(Patents) * S-track	12.47	6.09	10.89	9.68	-0.01	0.04	0.01	0.01
	(10.03)	(12.03)	(10.02)	(11.31)	(0.09)	(0.10)	(0.03)	(0.03)
Dot-gov * S-Track	51.06	125.70	55.06	144.53	0.09	0.54	0.30	0.11
	(77.17)	(95.42)	(84.70)	(93.70)	(0.92)	(1.05)	(0.17)*	(0.20)
Dot-org * S-Track	346.34	195.15	200.05	147.83	1.77	0.99	-0.39	-0.29
	(78.92)***	(102.35)*	(91.22)**	(104.58)	(0.86)**	(1.16)	(0.17)**	(0.20)
Dot-edu * S-Track	145.79	115.68	156.44	152.91	0.54	0.13	-0.22	-0.03
	(77.10)*	(82.16)	(70.23)**	(76.96)**	(0.54)	(0.61)	(0.15)	(0.16)
Suit-share	-2.44	-0.68			-0.03	-0.00	0.81	0.38
	(2.16)	(2.86)			(0.02)	(0.02)	(0.39)**	(0.59)
log(Patents)	-1.53	6.16	-4.95	8.99	0.05	0.07	-0.02	-0.03
	(8.29)	(11.55)	(7.79)	(10.97)	(0.07)	(0.09)	(0.02)	(0.03)
Dot-gov/org/edu	-89.28	-97.69	-53.14	-42.90	-0.67	-0.55	0.15	0.04
	(66.77)	(66.41)	(62.77)	(62.70)	(0.51)	(0.53)	(0.13)	(0.14)
ID Suit-share * S-track			4.38	4.11				
			(1.87)**	(1.91)**				
ID Suit-share			1.03	3.38				
			(1.48)	(1.68)**				
Tech Class Effects	Y		Y		Y		Y	
WG Effects		Y		Y		Y		Y
Large Firm Effects	Y	Y						
PubCohort Effects	Y	Y	Y	Y	Y	Y	Y	Y
Additional Controls†	Y	Y	Y	Y	Y	Y	Y	Y
Obs. (Standards)	671	671	671	671	671	671	926	926
Obs. (Nonstds)	287	287	287	287	287	287	482	482

*10% significance; **5% significance; ***1% significance (all SEs clustered on Working Group). **Notes**: Baseline sample contains all standards (and nonstandards) from Draft-Cohorts 1993 through 2002. All but Cox Hazard models omit observations with Total Days > 2007 (5.5 years) to correct for right-truncation of the DV, and use a propensity-score matched sample (see online appendix for details). †Additional controls are main effects and interactions for: log(Drafts), log(WG Mail), ID Sponsors, log(Filesize).

Table R-7: OLS Citation Models with WG Effects

Model	OLS (1)	OLS (2)	OLS (3)	OLS (4)	OLS (5)	OLS (6)
Dependent Variable	All Cites	All Cites	All Cites	RFC Cites	Patent Cites	Article Cites
Suit-share	0.39 (0.17)**			0.15 (0.06)**	0.16 (0.08)*	0.08 (0.04)*
Suit-share * S-track	-0.11 (0.09)			-0.08 (0.06)	-0.05 (0.03)*	0.02 (0.03)
log(Days)		1.18 (1.76)				
log(Days) * S-track		4.08 (2.33)*				
log(Versions)			0.78 (2.52)	0.83 (1.35)	1.18 (1.10)	-0.81 (1.01)
log(Versions) * S-track			7.48 (2.96)**	3.43 (1.52)**	1.69 (0.99)*	2.07 (1.28)
log(ID Mail)	3.76 (0.92)***	3.89 (0.91)***	3.34 (0.94)***	2.28 (0.58)***	0.32 (0.30)	0.38 (0.19)*
	Control Variables & Regression Statistics					
Draft Cohort FE's	10 [0.00]	10 [0.00]	10 [0.00]	10 [0.00]	10 [0.00]	10 [0.00]
RFC Month Polynomial	4 [0.00]	4 [0.00]	4 [0.00]	4 [0.00]	4 [0.00]	4 [0.00]
Additional Controls	12 [0.00]	12 [0.00]	12 [0.00]	12 [0.00]	12 [0.00]	12 [0.00]
WG Effects	186	186	186	186	186	186
Obs. (Standards)	872	872	872	872	872	872
Obs. (Nonstandards)	453	453	453	453	453	453
Model d.o.f.	29	29	29	31	31	31
Pseudo-LogL x 10^{-3}	-8.45	-8.35	-8.23	-6.39	-2.53	-2.28

*10% significance; **5% significance; ***1% significance (all SEs clustered on Working Group). **Notes**: The sample in models (1) through (6) is all RFCs from Draft Cohorts 1993 through 2003 published before 2007. For controls, the table indicates number of parameters and a p-value for joint significance. Additional controls are main effects and standards-track interactions for: S-track, log(Drafts), log(WG Mail), Sponsors, log(Filesize), Dot-edu, Dot-gov and Dot-org.