Standard Setting Committees: Consensus Governance for Shared Technology Platforms *

Timothy Simcoe

Boston University School of Management and NBER

June 11, 2010

*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: Boston University School of Management, 595 Commonwealth Avenue, Boston, MA, 02215. E-mail: tsimcoe@bu.edu

Standard Setting Committees: Consensus Governance for Shared Technology Platforms

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

Compatibility standards define key formats and interfaces for shared technology platforms. Network effects make standards self-enforcing, but hard to establish, since firms and users often delay adoption to see whether a new technology will succeed (Katz and Shapiro, 1985; Farrell and Saloner, 1986). Standard Setting Organizations (SSOs) help solve this coordination problem by providing a forum where interested parties can seek a broad consensus before endorsing a particular technology and promoting it as the industry standard.

In principle, SSOs promote orderly technical transitions that avoid the risk, duplication and intense competition of a decentralized standards war. But the switch from market to committee-based coordination need not alter participants' private interest in specific technologies. When distributional conflicts are strong, and participants lack the tools to craft an effective compromise, it can be hard to reach a consensus. Shapiro and Varian (1998, pg. 240) describe the formal standard setting process as "a wild mix of politics and economics." This paper asks when SSOs will work well.

I begin by developing a simple model of standard setting committees, based on the stochastic bargaining framework of Merlo and Wilson (1995). In this model, delays can be efficient, since they lead to better technological outcomes. However, these delays grow excessive when SSO participants favor specific technologies because of development lead times, proprietary complements, or intellectual property rights. The model predicts that coordination delays become longer and less efficient when the private benefits of adopting a preferred technology increase, or when the efficacy of side-payments and other concessions declines.

To test these predictions, I collect detailed committee and proposal-level data from the Internet Engineering Task Force (IETF), an influential SSO that produces standards used to run the Internet. The 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 those changes to construct a committee-level measure of commercial significance — the "suit-to-beard" ratio — that proxies for participants' private interest in specific technologies, and to create author-level proxies for the efficacy of side-payments.

To isolate 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-track RFCs" to publish new ideas without formally endorsing them. Since standards and nonstandards go through an identical publication process (often within the same committee) they are subject to the same unobserved sources of routine publication delay. Thus, I treat nonstandards-track RFCs as a no-conflict control sample and use them to estimate the relevant counterfactual: the time required to reach a consensus in the absence of distributional conflict.

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 7.8 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 eight months of deliberation for a typical standard. This effect has grown over time, and is larger for standards at the top of the protocol stack, where innovation in the underlying technology was more rapid. Delays are also 3 to 6 months longer for submissions from academic and non-profit authors, suggesting that IETF members from outside the private sector are less able (or willing) to make concessions.

Overall, the empirical results show that Internet commercialization caused an increase in strategic maneuvering within the IETF, and a slowdown in committee decision-making. These findings highlight the challenge of governing an economically significant piece of shared technical infrastructure. More, broadly, the results provide empirical support for the argument that rent-seeking is an important source of coordination costs. Many economists, such as Becker and Murphy (1992) and Jones (2008), cite coordination costs as a key factor limiting the gains from specialization. While these costs are often hard to measure, the IETF's transparency allows for a detailed examination of the link between rent-seeking and inefficient delays in technology adoption. Moreover, because compatibility standards are non-rivalrous and self-enforcing, SSOs provide a unique opportunity to isolate the link between rent-seeking and coordination delays in the absence of *ex post* monitoring and enforcement problems, which typically arise in the parallel problem of governing a shared congestible resource (Ostrom, 1990).

1.1 Related Literature

The literature on technical compatibility describes several paths to coordination. While competition is one possibility, standards wars can be intense and highly uncertain when network effects are strong (see, for example, David and Greenstein (1990) or Besen and Farrell (1994)). A second approach is for large "platform leaders" to orchestrate major technical transitions. But dominant platforms need not have dominant firms, as emphasized by Bresnahan and Greenstein (1999) in their historical account of the computer industry; and by Boudreau (2010), who highlights the distinction between platform access and ceding control over core technology. When there is no platform leader, a third path to inter-operability is for firms to create an institution for collective self-governance. In the information and communications technology sector, these groups are called Standard Setting Organizations.[1]

Since SSOs are voluntary organizations, and typically lack enforcement power, we might expect their recommendations to have little impact. But network effects can make cues for coordination self-enforcing. Rysman and Simcoe (2008) provide indirect evidence of this endorsement effect by documenting an increase in citations to standards-related 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.

Practitioners often point to delay as a source of 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]

Farrell and Saloner (1988) model delays in consensus standard setting as a war of attrition. Their theory predicts that SSOs are slower than markets, but more likely to achieve coordination. Farrell and 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. This paper models consensus standard setting as a complete-information stochastic bargaining game. Unlike a war of attrition, players can use concessions to reach agreement, and delays may be efficient, since the underlying technology improves over time. This new theory highlights two key inter-related drivers of SSO performance: the extent of distributional conflict, which creates incentives for rent-seeking, and the efficacy of side-payments, which provide a path to compromise.[4]

Lerner and 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 a proposed standard. Such forum-shopping suggests that distributional conflicts may be muted if firms with competing technologies join different SSOs. While this is an important insight, it is unlikely to influence the empirical results described below, since the IETF has emerged as the

[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, 2009).

[2] A substantial law and economics literature *assumes* that 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.

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

de facto forum for creating Internet infrastructure standards, with very few competitors.

The empirical literature on SSOs contains many descriptive studies (e.g. Besen and Johnson (1988); Besen and Saloner (1989); Farrell and Shapiro (1992); Foray (1995); Lehr (1995); Brown (1997); Levinson (2006)). While these qualitative accounts often describe distributional conflicts and link them to coordination delays, the rapidly growing quantitative empirical literature emphasizes a different set of questions. Several studies examine the link between SSOs and strategic alliances. For example, Rosenkopf et al. (2001) study alliance formation among firms that participate in the same standards committee. Bekkers et al. (2002) examines how SSO members' patent portfolios and alliance networks evolve over time. And Leiponen (2008) shows that a firm's position in an alliance network is correlated with the success of proposals to a focal SSO. Other empirical studies include Chiao et al. (2007), who analyze SSO bylaws and intellectual property policies; and Waguespack and Fleming (2009), who show that IETF participation is correlated with liquidity events in a sample of start-ups. This paper returns to the central normative question of the descriptive literature — is it better to coordinate through markets or committees? — and takes a first step towards understanding the drivers of SSO performance, by measuring the link between distributional conflicts among innovators and inefficient delays in the diffusion of new technologies.

The balance of the paper is organized into five parts. Section 2 presents a 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 Coordination Delays in Consensus Standardization

SSOs are governed by the consensus principle, which gives interested parties the power to block (or at least delay) the adoption of new standards.[5] The resulting process is typically modeled as a war of attrition, but that approach allows neither collaborative technology development nor multi-lateral bargaining. This section adapts the complete-information stochastic bargaining model of Merlo and Wilson (1995) to analyze coordination delays and the quality of outcomes in an SSO where technology improves over time and players use concessions to reach agreement. Equilibrium is characterized by a mix of efficient and inefficient delays that varies with the level of distributional conflict and the cost of using side-payments.

[5] The American National Standards Institute (ANSI, 2006) offers the following definition of consensus: "...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."

2.1 A Model of SSO Committees

Consider a symmetric committee with $k+1$ members. Each player sponsors a single proposal, and only one can be chosen as the industry standard.[6] Proposals vary in quality, which is indexed by a random variable $q \in [0, \infty)^{k+1}$ with symmetric joint distribution $F(\cdot)$ and continuous support. The element q_i is a publicly observed measure of the quality of player i's technology. In practice, q_i will depend on product market characteristics and aspects of the underlying design, such as technical performance, implementation cost, and flexibility.

At the start of each period, committee members pay a participation cost c, and receive an independent draw from $F(\cdot)$. If q_i improves on player i's last-period proposal, this draw becomes their new design, i.e. they are free to discard poor ideas. In the second half of each period, the committee member with the best technology makes a take-it-or-leave-it offer, which may include concessions $b \geq 0$, and the remaining players vote whether to accept it.[7] Consensus is modeled as a unanimity rule: standardization occurs if and only if all players accept a proposal. Upon rejection, all committee members receive an inside payoff, which I normalize to zero, and the game moves to the next period. The players' discount factor is $\beta < 1$.

If player i's offer is accepted, their present value payoff is $\Pi_i(q) = \pi(q_i, a_i; \omega) - kb$; 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 $\gamma \in [0, 1]$ reflects the efficacy of any concessions.

Since standards are public goods, I assume that $\pi(x, a; \omega)$ is increasing in x; everyone benefits from choosing a better technology. For simplicity, I assume the marginal benefits of improved quality are constant.[8] Winning also confers private benefits, since firms may profit by licensing their intellectual property or avoiding the costs of redesign. More broadly, winners can benefit from learning economies, proprietary complements, or time-to-market advantages if their preferred technology becomes an industry standard. To model the private benefits of winning, I assume that $\pi(x, a; \omega)$ is increasing in a, and has increasing differences in (a, ω), with $\pi(x, 1) = \pi(x, 0)$ when $\omega = 0$. These assumptions nest various 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.

[6] Large *ex ante* asymmetries will presumably lead to a swift resolution in favor of the player with better technology, more resources, or stronger incentives, as in Myatt (2005), suggesting the alternative institution of a predetermined platform leader. In practice, SSO committees with many participants are often characterized by a small number of roughly symmetric coalitions.

[7] Since this is an ultimatum-game, the player with the best technology has all of the bargaining power. While engineers often claim that technical merit plays an important role in SSO deliberations, a symmetric random-recognition rule would weaken the link between quality and bargaining power without altering any results. Allowing negative concessions ($b < 0$) also does not change the main results.

[8] This simplifies the statement and proof of Predictions 1 and 3, but is not a necessary condition.

Finally, committee members can use concessions b to "buy off" opposition to a proposal. Most SSOs 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.[9] 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 transaction costs of licensing.[10]

2.2 Equilibrium

In a symmetric Markov-perfect equilibrium, each player's strategy is a correspondence σ_i that maps the state variable q onto an offer and voting rule. If $p \equiv p(q, \sigma)$ is the equilibrium probability of adopting a standard, these strategies must solve the following program:

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

Since this model is stationary and symmetric, we have $V_{it} = V$. This implies a unique equilibrium at the voting stage: a proposal with quality x is accepted 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

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

Since a committee that is willing to adopt x would also accept any better proposal, these two inequalities define a reservation rule. The probability of achieving consensus is $p(x, \sigma) = 1[x \geq q^*]$, where q^* is the lowest-quality proposal that will be adopted in equilibrium. When $x = q^*$ the proposer's offer constraint (3) must bind, or they would be willing to make additional concessions to secure approval, thereby lowering the equilibrium quality threshold.

Replacing $p(x, \sigma)$ in Bellman's equation (1) yields V as a function of q^*. Plugging that expression into (2) and (3) and summing the system of equalities leads to the following con-

[9] 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.

[10] An alternative model that yields similar conclusions assumes that $\gamma = 1$, but players can only make (or credibly commit to) concessions below some threshold $\bar{b}$.

dition, where $G(\cdot)$ is the cumulative distribution of x (i.e. the first-order statistic of q), and $S(x) = k\pi(x, 0) + \pi(x, 1)$ is the gross surplus from adopting a standard:

$$\int_{q^*}^{\infty} \Big\{ S(x) - S(q^*) + k(\gamma - 1)[b(x) - b(q^*)] \Big\} dG(x) = (k+1)c + \frac{1-\beta}{\beta} \Big[S(q^*) + k(\gamma-1)b(q^*) \Big] \quad (4)$$

Equation (4) shows that the cut-off proposal quality q^* equates the marginal costs and benefits of continuation. The left side of (4) shows the expected benefits of search, given that q^* is the best-available proposal. The right side of (4) shows the search costs, which include the direct costs of participation and the opportunity cost of delayed implementation. Since the expected benefits of search are decreasing in q^*, while opportunity costs are increasing, there is a unique symmetric equilibrium.[11]

A realistic feature of this model, and a key challenge for empirical work, is that delays are sometimes efficient. For example, when concessions are as good as cash ($\gamma = 1$), equation (4) reduces to the optimal stopping rule for a single-agent search problem (e.g. Lippman and McCall, 1976), and the equilibrium cut-off rule maximizes the players' joint surplus. A committee will also adopt this first-best stopping policy if there is no distributional conflict ($\omega = 0$), and consequently no need for side-payments.

However, when consensus requires costly concessions, q^* will exceed the first-best. Intuitively, when $\omega > 0$ losers wish to bargain beyond the optimal stopping threshold, since it could lead to the adoption of their preferred technology (or equivalently, an increase in bargaining power). When $\gamma < 1$, a proposer with quality q^* is unwilling to offer concessions that fully compensate losers for the loss in bargaining power, since that would lead to a payoff below their continuation value. Thus, the search for consensus takes too long (in expectation) if and only if there is both distributional conflict and costly concessions.

2.3 Testable Implications

I test three predictions of this model, starting with the impact of a change in distributional conflict. Increasing ω widens the gap between winning and losing payoffs, which leads to more rent-seeking as players hold out for their own technology. It also changes the gross payoff to standardization $S(x)$, which alters the opportunity cost of delay. Prediction 1 provides sufficient conditions for increased conflict to produce longer (expected) delays.

Prediction 1: Coordination delays (q^*) are weakly increasing in ω if opportunity costs also decline ($S_\omega(x) \leq 0$); or if players are sufficiently patient ($\beta = 1$); or if there is no redistribution

[11] See the appendix for a proof of all results discussed in the remainder of this section.

($\pi_\omega(x,0)=0$) and concessions are sufficiently costly ($\gamma=0$).

Intuitively, when $S_\omega(x) \leq 0$ increased rent-seeking and decreased opportunity costs both promote longer delays. This condition would hold for any purely redistributive change in payoffs, such as an increase in fixed royalty payments that losers cannot pass on to customers. It would also hold in a model of cost avoidance, where losers must redesign their products to conform with the winner's favored technology, and those costs increase with ω.

When $S_\omega(x) > 0$ rent-seeking and the increased opportunity costs of delay work against one another. However, opportunity costs are unimportant when players are patient (i.e. as $\beta \to 1$). And even when β is small, the rent seeking effect will dominate if concessions are very costly (i.e. as $\gamma \to 1$) and increasing ω primarily helps the winner; as in a model where inelastic final demand allows losers to pass increased royalty payments on to consumers.

Prediction 2: Coordination delays (q^*) are weakly decreasing in γ.

Less costly concessions make the consensus process more efficient. This result holds under very general conditions because γ has no impact on gross payoffs or opportunity costs.

Prediction 3: If coordination delays (q^*) increase with ω, the rate of change is weakly decreasing in γ, so that $\frac{d^2 q^*}{d\omega d\gamma} \leq 0$.

Intuitively, players can respond to increased distributional conflict in two ways; by raising q^* or offering more concessions. As γ declines, concessions become more costly and they naturally shift toward using the cut-off rule.

Since all three predictions work through a change in q^*, it is tempting to conclude that a decline in the efficacy of concessions or an increase in distributional conflict will lead to better standards, albeit more slowly. But this depends on how one interprets the cost of concessions. If $\gamma < 1$ because of exogenous factors, unrelated to technical quality, which limit the use of monetary side-payments, then distributional conflict and *ex post* technical quality should be positively correlated. However, if γ represents technical inefficiencies that arise through design by committee, such as efforts to satisfy all players' potentially conflicting demands, the positive association between conflict and expected quality could be quite weak.

I now turn to the empirical analysis, which seeks to isolate the rent-seeking effects highlighted by this model in the context of Internet standards development.[12]

[12]The model generates other comparative statics that I do not test. For example, increasing β leads to longer delays. While increasing k has an ambiguous impact on q^*, one can show that q^* approaches the first-best as k grows large, since rent-seeking incentives are proportional to $\frac{1}{k+1}$, the *ex ante* probability of winning.

3 The Internet Engineering Task Force

The IETF creates and maintains the standards used to run the Internet, such as the Transmission Control Protocol and Internet Protocol (or TCP/IP) for routing packets, and DHCP for assigning network addresses. The organization was formed in 1986, and early members were primarily academic and government researchers.[13] Commercial interest was limited by the NSF's Acceptable Use Policy, which prohibited commercial Internet use before 1991, and by the availability of competing standards, such as IBM's Systems Network Architecture.

During 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.[14] Table 1 shows the increase in IETF committees and proposals during this period. There was also a significant shift in member demographics, as individuals affiliated with private-sector organizations grew from roughly 50 percent of IETF participants in 1993 to over 80 percent by 2001.[15]

The growth and commercialization of the IETF was driven by increased demand for new protocols to extend the functionality of the Internet. For example, new applications such as instant messaging and voice-over-IP required new standards for user authentication, security, and session management. Much of this activity occurred near the top the TCP/IP protocol stack, a conceptual model of a communications network with five hierarchical layers: Application, Transport, Internet, Link, and Physical (or Hardware). The stack is an integral part of the Internet's end-to-end architecture (Saltzer et al., 1984), which relies on low-level protocols to provide fast reliable packet delivery, and leaves more complex functions to higher-level protocols near the "edges" of the network. This division of labor among complementary standards recognizes that adding features at the core of a network often incurs costs at higher-layers, and pushes the locus of innovation — where standardization is most likely to impinge on proprietary technology — away from that core.

While rapid innovation at the top of the stack was partly a consequence of TCP/IP becoming a *de facto* standard, it also had important consequences for Internet governance. In particular, the IETF became a place where vendors could seek to ensure that proprietary technology would work within the global Internet. Anecdotal evidence suggests that growth and commercialization led to increased tension within the IETF (Davies, 2004). For example, a committee working on instant messaging protocols received proposals from both Microsoft and

[13] Meetings in 1990 drew about 100 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).

[14] Detailed histories of this period include Russell (2006), Abbate (1999) and Berners-Lee and Fischetti (1999).

[15] These figures are based on the top-level domain of emails sent to any IETF Working group listserv. The appendix provides additional descriptive statistics to illustrate IETF growth and commercialization.

AOL, who were in the midst of a standards battle over the issue. In another committee, Cisco clashed with a pair of start-ups over standards for wireless access points.[16] Brim (2004) describes several cases where firms claimed intellectual property in technologies being evaluated by the IETF. This paper measures the impact of increased distributional conflict on the pace of IETF standards development.

3.1 Standard Setting Process

The IETF has its own language for describing the standard setting process. Figure 1 provides an overview. Committees are called Working Groups (WGs), and proposals are called Internet Drafts (IDs). A published ID is called a Request for Comments (RFC). There are two types of RFC. Standards-track RFCs define new protocols, which progress in maturity from Proposed Standard to Draft Standard to Internet Standard. Nonstandards-track RFCs are classified as Informational or Experimental.[17] While standards and nonstandards go through an identical publication process, as described in Bradner (1996), nonstandards do not receive an official endorsement, and may advance no further unless re-subitted as an ID for standards-track publication.

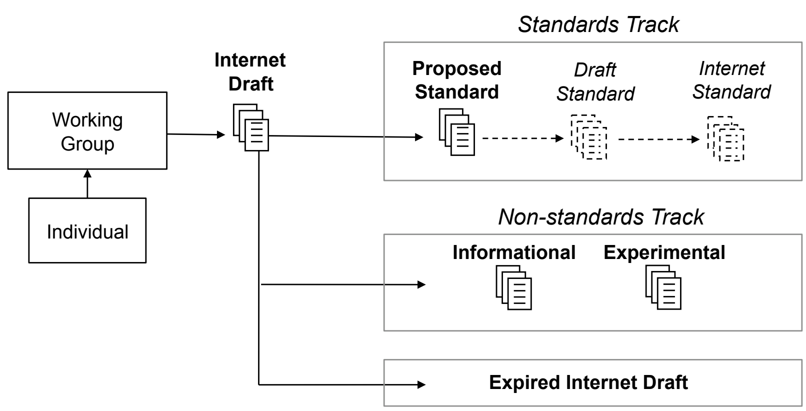

Figure 1: The IETF Standard Setting Process

The IETF standards development process begins when participants identify a problem and form a Working Group to consider solutions. To prevent forum shopping and overlapping technical agendas, new WGs must be approved by an advisory board called the Internet Engi-

[16] "WLAN standards battle begins again," Marguerite Reardon, CNET News, April 11, 2005.
[17] I ignore a third type of nonstandards-track RFC, called a Best Current Practice (BCP), that is used to describe IETF policies and procedures.

neering Steering Group (IESG).[18] Once a WG is formed, anyone can submit an Internet Draft by posting it to a public repository.[19] New IDs are debated at tri-annual IETF meetings, and on the e-mail discussion lists (or listservs) maintained by each WG, where much of the organization's work is done. Internet Drafts are continually revised, and an unpublished ID will expire after six months if the authors do not submit a revision.

For an Internet Draft to become an RFC, the relevant WG must reach a "rough consensus" on the merits of the proposal. While the IETF provides no formal definition, rough consensus is often described as the "dominant view" of the Working Group, and implies support from well over 51 percent of active participants (Bradner, 1998, pg. 12). In practice, a WG chair decides whether consensus has emerged. If the WG chair declares a consensus, there is a "last call" for comments within the WG, and the ID is submitted to the IESG. The IESG reviews the proposal and issues a second last call for comments from the entire IETF community. Any comments or formal appeals are reviewed by the IESG, and may be referred back to the WG for resolution. If the IESG is satisfied that a consensus exists within the WG, and sees no problem with the ID, it will be published as an RFC.

3.2 Standards and Nonstandards

The distinction between standards and nonstandards-track RFCs plays a central role in the empirical analysis, where I use nonstandards to control for omitted variables that may be correlated with both distributional conflict and standards-track publication delays. The key difference between the two publication tracks is that standards-track RFCs receive the IETF's formal endorsement, while nonstandards do not. Thus, while standards provide a commercially relevant focal point for implementation, which can produce winners and losers, there is no comparable incentive to prevent or delay the publication of nonstandards.

Standards are a prescriptive signal to technology developers that, "if you are going to do something like this, you must do exactly this" (Postel, 1995, pg. 10). To assess compliance with standards-track RFCs, the IETF defines a set of standards-track keywords, which authors use to assign a formal requirement-level to each feature in a new protocol.[20] Nonstandards-track RFCs are descriptive. Instead of providing a focal point for widespread implementation, nonstandards offer a convenient outlet for publishing useful information. While standards-track keywords may appear in a nonstandards-track RFC, they would have no formal meaning, since

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

[19] Individuals can submit IDs outside of the WG process. Since these individual IDs are unlikely to become an RFC, and cannot become a standards-track RFC, I exclude them from the empirical analysis.

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

there is no obligation to "comply" with an Informational or Experimental RFC.

Working Groups use the nonstandards-track in two ways. First, a nonstandards-track RFC may describe ideas that are too preliminary or controversial to become a standard. For example, the Experimental RFC 2582 suggests changes to TCP to help manage network congestion. While the IETF did not initially endorse the proposal, it was published as a nonstandard to encourage further experimentation, and the underlying ideas were later re-submitted for standards-track publication. The second use of nonstandards is to provide information that complements a standard, such as guidelines for implementation and deployment. For example, Informational RFCs have been used to catalog the negative externalities that occur when vendors fail to comply with a protocol (RFC 2525), and to propose a network architecture based on protocols defined in a set of related standards (RFC 2475).

As described below, the empirical analysis uses nonstandards-track RFCs as a no-conflict control sample. This approach makes two key assumptions. First, I assume that standards and nonstandards — which go through an identical publication process, often within the same WG — suffer the same routine publication delays, such as a slow WG chair, a complex technology, or congestion in the review process. And second, I assume that nonstandards produce no distributional conflicts, so there is no correlation between conflict and delay.

The main weakness of the nonstandards-track control sample is that an ID's publication track is not observed until the last call announcement (which never occurs for expired IDs). This raises the possibility that standards could morph into nonstandards while under review, or vice versa. In practice, IETF participants typically know whether an ID is on the standards or nonstandards-track after one or two revisions, and suggest that expired proposals are often failed standards, since there is no reason to reject a nonstandards-track proposal that passes some minimum quality threshold.[21] Nevertheless, the next section outlines an identification strategy that uses standards-track keywords as an instrumental variable to control for endogenous selection onto a particular track, based on the idea that keywords measure an ID's intended publication-track when it is first submitted.

4 Empirical Strategy

This section develops an empirical strategy for measuring the link between commercialization and delays in the IETF standard-setting process, and describes the data used in estimation.

[21] There are exceptions, such as the "Link-Local Multicast Name Resolution Protocol" (RFC 4795), which went through 48 standards-track revisions before publication as an Informational RFC.

4.1 Empirical Model and Identification

Suppose we observe a cross-section of proposals i, submitted to committees j at times t. These data include the time-to-consensus T_i for each proposal, a vector of proxies D_{ij} for distributional conflict or the efficacy of side-payments, an indicator I_i that equals one for standards-track and zero for nonstandards-track RFCs, and a vector X_{ijt} of additional controls that includes a constant and a full set of time-period and technology-class dummies.

To measure the impact of distributional conflict on coordination delays, one could regress T_i on D_{ij} in the sample of standards-track RFCs. Unfortunately, this approach is vulnerable to omitted variables. For example, if D_{ij} is correlated with delays caused by unobserved technical complexity, a linear regression would overstate the rent-seeking effect. In a controlled setting, one might address the omitted variables problem by randomly selecting proposals 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.

As described above, standards and nonstandards experience the same routine publication delays, since they focus on similar technology and go through an identical publication process. However, the commercial stakes are lower for nonstandards, which are not meant to provide an impetus for coordinated implementation. I assume there is no rent-seeking on the nonstandards-track — which implies no correlation between rent-seeking and delays — and use the nonstandards to estimate a counterfactual correlation between D_{ij} and unobserved sources of routine publication delay.[22] This leads to the familiar difference in differences estimator

$$T_i = I_i\alpha + D_{ij}\beta + D_{ij}I_i\tau + X_{ijt}\theta + \varepsilon_i \tag{5}$$

where α measures the mean difference in time-to-consensus, e.g. because standards are more complex; β measures the correlation between D_{ij} and routine publication delays; and τ measures the impact of distributional conflict on delayed agreement. When D_{ij} is continuous, τ is estimated by a difference in the slope of two regression lines, as opposed to a difference in intercepts.

This diff-in-difs framework will produce unbiased estimates when I_i is exogenous or τ is constant. In practice, the IETF does not randomly assign proposals to a particular track, so estimates based on (5) are vulnerable to selection bias when there are heterogeneous effects. For instance, authors may steer proposals that encounter unexpectedly fierce resistance onto

[22] One might argue that zero is a theoretical lower bound on distributional conflict, since people can always find something to fight over. However, a positive correlation between conflict and delay for nonstandards would bias estimates of τ downwards. In practice, IETF participants suggest that nonstandards receive little opposition.

the nonstandards-track, forgoing the certification benefits of becoming a standard in favor of a less costly review process. Such morphing of standards into nonstandards would produce a downward bias if more conflict (higher τ) leads to selection out of the standards-track treatment group. However, the sign of this bias could be reversed if unobserved certification benefits are highly correlated with distributional conflict. I use several methods to guard against either possibility.

First, I use propensity-score matching (Rosenbaum and 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. in technical complexity. Third, I use standards-track keywords to instrument for I_i in the diff-in-diffs model. And fourth, I use the same instrument to estimate a switching model that treats standards and nonstandards as separate regimes — each with its own duration equation — and adds the following selection equation:

$$I_i = 1[Z_i\pi + \nu > 0] \tag{6}$$

A parametric switching model assumes that the unobservables in (6) and the two duration equations have a trivariate-normal distribution: $(\varepsilon^s, \varepsilon^n, \nu) \sim N(0, \Sigma)$, where s, n index standards and nonstandards respectively; 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. This model was introduced by Heckman (1976), and Vytlacil (2002) shows that if one drops the distributional assumptions, it identifies the same Local Average Treatment Effect as a linear IV analysis. I include the switching model as a complement to the instrumental variables estimates. While the switching model makes stronger distributional assumptions, it produces more stable estimates when $D_{ij}I_i$ contains many endogenous interaction terms, and (as described below) provides a simple framework for incorporating censored and expired IDs.

In both IV regressions and the switching model, I use a count of standards-track keywords contained in the first version of an ID to instrument for I_i. Intuitively, standards-track keywords measure an ID's "intended" publication-track. Thus, instrumenting for I_i will control for standards morphing into nonstandards (or vice versa) based on information that arrives during the review process. This IV strategy may break down if there is an *ex ante* selection process that precedes the creation of a first draft. In that case, an IDs intended track might be a function of expected delay, as opposed to an exogenous "idea-generating process" that determines standards-track suitability. The IV would also fail if keywords are correlated with unobserved specificity, scope or other factors that contribute to coordination delays.

Within the switching model, standards-track keywords can also be used to incorporate censored and expired IDs into the analysis. The baseline diff-in-diffs model drops these observations (about 40 percent of the total sample) because I_i and T_i are not observed for unpublished IDs. However, the selection equation (6) uses the correlation between Z_i and (observed) I_i to estimate of the probability that an ID is on a particular track. Assuming all expired proposals would eventually become an RFC — so they are merely censored as of the final submission date — predicted values from the switching model can be used to compute the total likelihood for a given observation, which is just $Pr(I_i = 1, T_i > T) + Pr(I_i = 0, T_i > T)$. This leads to a switching model with "partially observed regimes" which I estimate via maximum likelihood.[23]

4.2 Data and Measurement

All data come from the IETF's public archives, and are available through the author's web site.[24] The population consists of 3,521 Internet Drafts submitted to IETF Working Groups between January 1993 and December 2003, and the estimation sample contains 2,601 IDs that went through at least one revision.[25] 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.

I measure coordination delays as time-to-consensus, starting with the submission of an ID, and culminating in one of three ways: publication as a Proposed Standard, publication as a nonstandards-track RFC, or expiration.[26] 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. 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).

Table 1 shows mean delays by publication-type and submission-year. The average time-to-consensus for a Proposed Standard (standards-track RFC) was 774 days (2.1 years), compared

[23]See the appendix for a detailed description of the likelihood function. An alternative approach assumes that all nonstandards-track IDs are published, so expired proposals must be on the standards-track. I explore this idea using a Cox hazard model in the robustness checks.

[24]All data and code, plus more details on variable creation, can be found at http://people.bu.edu/tsimcoe.

[25]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

[26]I omit Draft Standards and Internet Standards from the analysis, since they are not new protocols, but rather a formal recognition of widespread implementation and deployment.

to 595 days for a nonstandards-track RFC and 487 days for an expired proposal. While time-to-consensus is clearly increasing on both tracks, linear regression reveals a statistically significant difference in time-trends, which increase by 49 days per-year on the standards-track and 31 days per-year on the nonstandards-track. Table 1 shows that Proposed Standards submitted in 1997 and 1998 experience a sharp drop in publication lags relative to nonstandards. However, I estimate models that allow the rent-seeking effect to vary by Draft-Cohort and find no evidence of a systematic decline during this period.[27]

Table 2 provides a short definition and summary statistics for all variables used in the analysis. 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." Suit-share is defined as the percentage of all email domains (e.g. ibm.com) from dot-com or dot-net organizations on a WG listserv during a one-year window prior to the initial submission of an ID.[28] The idea behind this measure is that distributional conflict increases 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. Higher Suit-share is a good indication that the Working Group is creating commercially relevant technology. While Suit-share may also reflect the size of the (perceived) market opportunity, the model in Section 2 suggests that without conflict, larger payoffs lead to faster decisions given the increased opportunity cost of delay. In practice, distributional conflict often emerges when firms recognize an opportunity and race to enter 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).[29] 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 WG in a given year. While this second proxy for distributional

[27] One former IESG member suggests that this might reflect the urgency surrounding standards that were crucial to large network operators' expansion plans, or the "fast tracking" of web-related protocols such as RADIUS and SSL (private email communication).

[28] 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 the appendix.

[29] 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 remove some the largest firms, such as IBM, and find that it makes no difference.

conflict is closely linked to a particular mechanism (licensing), it is quite noisy because of the skewed size distribution of firms' patent portfolios and the difficulty of identifying the handful of patents essential to a particular standard.

The third panel in Table 2 contains a series of dummy variables that identify proposals from the academic and non-profit research communities. Specifically, Dot-org and Dot-edu equal one if an ID has one or more authors from the corresponding top-level domain, and Org-edu equals one if either Dot-org or Dot-edu is positive.[30] I use these indicators as a proxy for less effective concessions (i.e. smaller γ). 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 make concessions that appeal to large technology vendors, since they have different objectives and 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 Matching

The baseline diff-in-diffs specification of equation (5) 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 standards and nonstandards. While these means are indeed quite similar, I use propensity-score matching to create a sample where they are balanced by construction.[31]

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 longer to publish, go through more revisions, are mentioned in more emails, and contain more standards-track keywords than nonstandards-track RFCs. Standards also receive more forward-citations from RFCs and U.S. patents.[32] All of these differences are consistent with the maintained assumption that standards-track RFCs have

[30]Although participants have a wide range of affiliations, most Dot-edu contributors come from tier-one research universities, while Dot-org contributors are frequently from professional societies, related technical consortia and research institutes (e.g. W3C, Korea's ETRI, or the Internet Mail Consortium).

[31]Specifically, I trim the sample by dropping RFCs whose probability of becoming a Proposed Standard (based on fitted values from a probit) is below the 5th percentile of the empirical distribution for Proposed Standards, or above the 95th percentile for nonstandards-track RFCs. See the appendix for details.

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

greater commercial impact, and therefore generate more distributional conflict. Interestingly, standards and nonstandards are cited at roughly the same rate by academic journal articles, suggesting that academic authors are less concerned with the commercial implications of the underlying idea (though journal citation rates are also quite low).

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 proxies for distributional conflict and the efficacy of concessions, the only meaningful difference appears on Dot-edu, indicating that academic authors 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 that have evaluated more IDs and generated more email traffic. Though small, these differences are statistically significant, and Hotelling's T^2 test rejects the hypothesis that the sample means of all variables are equal. For the matched sample, however, standards and nonstandards are statistically indistinguishable. While the age and cumulative output variables still exhibit small differences, the T^2 test adjusts for the fact that these variables are highly correlated. 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. The main finding is that commercialization, as measured by Suit-share, is positively correlated with an increasing difference between standards and nonstandards-track delays.

5.1 Conflict, Concession and Coordination Delays

Figure 2 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 nonstandards-track delays exhibit little or no change. This finding is consistent with Prediction 1: delays increase with distributional conflict. To control for time-trends, technology effects or other factors that might influence this pattern, I turn to the difference-in-differences specification in equation (5), which assumes that selection onto the standards-track is exogenous, and measures the rent seeking effect by interacting measures of distributional conflict with a standards-track dummy variable. The results are presented in Table 4.

Column (1) presents OLS diff-in-diff 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

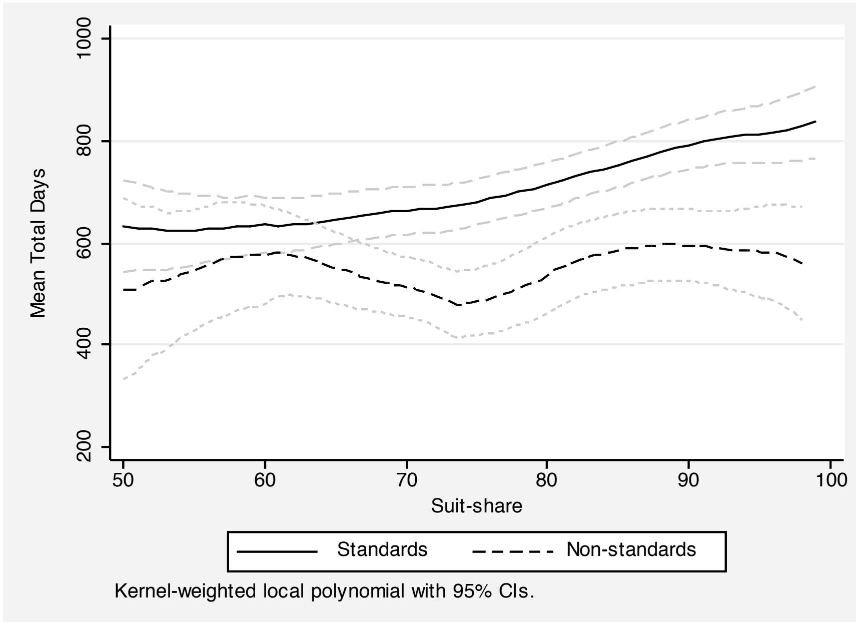

Figure 2: Difference in Duration by Suit-Share

WG fixed effects, identification comes from variation across RFCs both within and between committees. The interaction between Suit-share and S-Track in model (1) is large and statistically significant: a one percentage-point increase in Suit-share adds 7.8 days to the publication process. The interaction of log(Patents) with S-Track is smaller and statistically insignificant, suggesting that a one standard deviation increase in background IPR increases delay by roughly 20 days.[33]

Model (1) also includes the Org-edu indicator, which proxies for the cost of concessions. The coefficient on the interaction between S-Track and Org-edu implies that Dot-org and Dot-edu authors have a greater difference between standards-track and nonstandards-track publication delays than authors from a commercial top-level domain. This effect is large (just over 5 months) and consistent with Prediction 2: delays grow longer as the efficacy of concessions declines.[34]

[33] In general, I find that log(Patents) has a statistically significant positive correlation with standards-track delays (see model (5)), but is not strongly correlated with the difference in time-to-consensus across tracks. This might reflect attenuation bias, a positive correlation between the Suit-share and log(Patents) proxy variables, or an omitted variables problem that is addressed through the nonstandards-track controls.

[34] If I allow separate main and interaction effects for Dot-org and Dot-edu, I find a larger difference between standards and nonstandards-track delays for academic authors (162 versus 123 days). While there is no evidence of a difference in delay across tracks for Dot-gov authors, the results change very little if I add them to the Org-edu dummy (the point estimate on the Org-edu-gov interaction effect is 124 days (p=0.06)).

Column (2) in Table 4 restricts attention to the the matched sample and finds results that are statistically indistinguishable from model (1). The only notable change is a small decline in the Suit-share interaction effect (which falls to 6.8 days per Suit-share point) caused by an increase in the nonstandards-track Suit-share coefficient from -3.3 to -1.3.

Model (3) adds Working Group fixed-effects to control for time-invariant committee-level heterogeneity, so identification comes from variation among RFCs produced by a single IETF committee. The main results all fall within the 95 percent confidence intervals for model (1). However, the Suit-share interaction coefficient drops to 5.4, and the Org-edu effect declines to 92 days, which is not statistically significant.[35] While model (3) shows that the link between Suit-share and standards-track delay is robust to WG level unobserved heterogeneity, fixed-effects inference may be overly conservative. In particular, a Hausman test fails to reject the consistency of random effects estimates that are very similar to column (2), with a Suit-share interaction effect of 6.7 (p=0.00) and an Org-edu interaction coefficient of 136 (p=0.06).

Models (1) through (3) provide strong evidence of a causal link between IETF commercialization and coordination delay. They also document a larger difference between standards- and nonstandards-track publication delays for academic and non-profit authors, perhaps due to a comparative disadvantage at offering concessions. However, all of these results assume exogenous selection onto a particular publication track. Model (4) uses standards-track keywords to address the possibility of selection bias. Specifically, I estimate a probit selection model, with Keywords and log(KeyCount) included in Z_i, and use fitted values from the selection model as instruments for S-track in a linear IV specification. For the interaction terms, I construct instruments by interacting the probit fitted values with the exogenous regressors. Wooldridge (2002, pg. 626) describes this estimation strategy, and notes that fitted values from equation (6) will be optimal IVs if the selection model is correctly specified. Given previous results, I drop the log(Patents) interaction term to keep the number of endogenous variables small.

Column (4) in Table 4 presents results from this "two step" IV estimator. The Suit-share interaction coefficient is 12.3 (p=0.04), while the Org-edu effect is 264 (p=0.12).[36] Both of these coefficients are roughly 50 percent larger than in the model with exogenous selection, though a chi-squared test fails to reject the null hypothesis of exogeneity. One interpretation of the larger IV effects is that OLS estimates of the rent-seeking effect are biased downward.

[35] Allowing separate Dot-org and Dot-edu interaction effects in (3) yields a Dot-edu coefficient of 185 days (p=0.02), and an insignificant Dot-org effect of -26 days (though very few WGs produce multiple nonstandards where one or more has a Dot-org author (see Table A4)).

[36] While the estimates in (4) are just identified, it is possible to obtain overidentifying restrictions by interacting the probit fitted values with additional exogenous regressors (assuming those variables have an identical main effect for standards and nonstandards). In these overidentified models, Hansen's J statistic never rejects the hypothesis of a valid set of instruments.

In particular, much of the difference between models (1) and (4) comes from a decline in nonstandards-track Suit-share and Org-edu coefficients, as we might expect if especially controversial standards-track IDs were morphing into nonstandards. A second possibility is that I have weak instruments. However, keywords are highly significant in the first-stage probit, whose marginal effects show that a single Keyword increases the probability of becoming a standard by 5 percent, and that each unit of log(KeyCount) raises that probability by an additional 3.2 percent. Table 4 also reports the Angrist and Pischke (2009, pp. 217-18) first-stage F statistics for each endogenous regressor, which are well above the Stock and Yogo (2005) critical thresholds for high levels of bias. Overall, model (4) suggests that the baseline diff-in-diff results are, at worst, conservative; thus strengthening the argument for a causal link between IETF commercialization, distributional conflict and coordination delays.

The final column in Table 4 drops the nonstandards-track baseline and estimates a random-effects model for the standards-track sub-sample. In this specification, all of the proxies for distributional conflict and the cost of concessions are positively correlated with publication delays. The Suit-share effect declines to 4.2 days per percentage point, but remains statistically significant at the 1 percent level. The log(Patents) coefficient increases in both size and significance, while the Org-edu effect falls to 87 days. The main message of model (5) is that the baseline diff-in-diff results are not driven entirely by the nonstandards-track controls, and that background IPR is positively correlated with standards-track delays.[37]

The main 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 right-censored RFCs that have especially long publication lags. Using one-year lags of Suit-share as an instrument for Suit-share produces slightly larger estimates of the rent-seeking effect. Replacing the WG-level Suit-share variable with an ID-level Suit-share variable (based on emails that mention a specific ID while it is under evaluation) also produces similar results. Quantile regressions on the standards-track sub-sample show that delays at the median and 75th percentile increase by 4.0 and 6.6 days per Suit-share point respectively, while nonstandards-track quantile regressions show no relationship between Suit-share and delays. Finally, if one aggregates these data to the WG-Year-Track level and estimates a panel model with WG fixed-effects, the Suit-share interaction remains statistically significant.[38]

[37]A fixed-effects model for the standards-track sub-sample produces similar estimates, as evidenced by the Hausman test, which fails to reject the null hypothesis that WG effects are orthogonal to observables.

[38]See appendix Tables A5 and A6 for additional detail and estimates from several of these robustness checks.

5.2 Switching Model

Table 5 presents estimates from the switching model described in Section 4.1. Model (1) uses the Full Sample of standards and nonstandards-track RFCs (as in Table 4). Model (2) includes all censored and expired IDs. I estimate both models via maximum likelihood, with Log(Total Days) as the dependent variable.[39] Because of the log transform, the duration equation coefficients can be interpreted as the elasticity of the constant hazard rate for publication.

The row labelled "Weak IV's" in the bottom panel of Table 5 shows that both models strongly reject the null hypothesis that the Keyword coefficients are jointly zero. Interestingly, unreported OLS regressions show that the Keywords indicator is uncorrelated with standards-track time-to-consensus, but positively correlated with nonstandards-track delays. That finding is consistent with claims by IETF participants that standards-track IDs occasionally morph into nonstandards, while the opposite transition does not occur.

Turning to the estimates, model (1) shows that controlling for endogenous selection onto the standards-track does not change the baseline Suit-share result. In particular, Suit-share has a positive and highly significant coefficient in the standards-track duration equation, a negative coefficient in the nonstandards-track duration equation, and a Wald test strongly rejects the hypothesis that the two coefficients are equal. While the two Org-edu coefficients are statistically insignificant, the difference between them indicates that the gap between standards and nonstandards-track delays is 20 percent larger for Org-edu authors, and a Wald test finds this difference significant at the 10 percent level. Finally, since the switching model produces estimates of the correlation between unobserved error-terms across equations, it is possible to test whether $\sigma_{s\nu}$ and $\sigma_{n\nu}$ are jointly equal to zero. The last row in Table 5 shows that for the full sample of published RFCs, I cannot reject the null hypothesis of exogenous selection.

Model (2) in Table 5 presents estimates from a switching model that exploits information from all IDs submitted to the IETF between 1993 and 2003, including censored and expired proposals. Once again, the difference in Suit-share coefficients is large and statistically significant. While the Org-edu coefficients are roughly the same size as in model (1), the difference between them is no longer significant. And once again, model (2) fails to reject the null hypothesis that selection is exogenous.

5.3 Treatment Interactions

This sub-section examines heterogeneity in the link between distributional conflict and coordination delays. Since controlling for endogenous selection did not alter the main Suit-share

[39] Using a two-step Heckit procedure to estimate the switching model in levels yields results very similar to those presented in Table 4. That approach also fails to reject the null of exogenous selection.

results, I return to the simpler OLS diff-in-diffs framework, but allow the rent-seeking effect to vary with observables. This leads to the following specification

$$T_i = \lambda_k + I_i\alpha_k + D_j\beta_k + D_jI_i\tau_k + X_{ijt}\theta_+\varepsilon_i \tag{7}$$

where k indexes three types of heterogeneity: the Dot-org and Dot-edu indicators (to test Prediction 3), a set of IETF technology class dummies, or a set of time-period effects. In each case, I estimate two models. The first model uses only standards-track RFCs to see whether there is any evidence of a treatment interaction within the treated population. The second specification is the triple-difference model of equation (7), which includes main effects and two-way interactions for S-track, Suit-share and the source of heterogeneity, along with the three-way interaction of interest (τ_k). The results are presented in Table 6.

Models (1) and (2) interact Suit-share with Dot-org and Dot-edu, and find a significantly larger correlation between Suit-share and standards-track delays for proposals with a Dot-org author. This result is consistent with Prediction 3: the marginal impact of increased distributional conflict increases with the cost of concessions. While the Dot-org result is significant at the 5 percent level in the standards-track sub-sample, and at the 10 percent level in the triple-difference specification, there is no evidence of an interaction effect for Dot-edu authors. Unreported treatment interactions for the combined Org-edu indicator variable are positive and significant at the 10 percent level in either specification.

The bottom panel of Table 6 reports a Wald test of the hypothesis that the two-way interactions between Suit-share and author-type are jointly zero. This can be viewed as a specification test for the nonstandards-track control sample, since the three-way interaction effects will be identical to the more precise estimates in column (1) if those two-way interactions are omitted. For model (2), the Wald test does reject the null that the two-way interactions are jointly zero, which is not surprising given the small difference in estimated coefficients across (1) and (2).

Models (3) and (4) 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, with the Applications Area at the top, and the Routing Area at the bottom (Security and Operations do not occupy a specific location within the stack). The two models produce similar results, with estimates of the rent-seeking effect varying from -1.1 to 11.7 days per Suit-share point across the six technology areas. Both models find more rent-seeking at the top of the protocol stack, in the Applications and Transport areas. This finding is consistent with earlier discussion of the Internet's end-to-end design principles: low level protocols, such as IP, were already established by the time of rapid IETF commercialization, and remained relatively stable as higher level protocols were built on top of them. Rent-seeking was most pronounced at higher levels because that was the locus of innovation, where new applications created demand for new protocols which threatened the profits (or sunk investments) associated with proprietary solutions. Thus, the results in models (2) and (3) reinforce this paper's broader message by showing that the impacts of distributional conflicts are most severe when opportunities for private exploitation are high relative to losses from failing to solve a collective-action problem.

Finally, Models (5) and (6) examine variation in the rent-seeking effect over time. 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. Model (5) shows a drop in the correlation between Suit-share and delays for IDs submitted in 1997 and 1998, corresponding to the drop in average delays seen in Table 1. Model (6) allows the nonstandards-track Suit-share parameter to vary by Draft Cohort, and finds a steady increase in rent-seeking over time. While I cannot reject the hypothesis that the nonstandards-track interaction effects in model (6) are jointly zero, comparing the two models suggest that there was an acceleration on both tracks in during the late 1990s. The "double dip" in delays for those Draft Cohorts might reflect IETF policy changes or technological vintage effects (possibly associated with a profusion of web-related protocols), but it did not alter the underlying trend in the relationship between commercialization, rent-seeking and coordination delays.

5.4 Discussion of Results

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 understand the magnitude of this effect, consider the 30 percentage point increase in Suit-share across all Working Groups between 1993 and 2003. If each one-point increase led to 7.8 days of deliberation, then commercialization added roughly 7.8 months to the length of the IETF standard setting process. This effect would be economically significant in many Internet markets, 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 the IETF successfully managed a very challenging transition from research-oriented to commercially relevant standards organization.

While these results clearly link distributional conflict to SSO performance, they do not have straightforward welfare implications. For example, I do not quantify the costs or benefits of timely standard setting for end-users, for whom long delays might even be efficient if it leads to a smoother migration path. Welfare effects also depend on the technical quality of standards, which can be difficult to measure, and on the costs and benefits of alternative institutional arrangements, such as a standards war or platform leader.[40] Nevertheless, the consensus within the standard setting community is that rapid technological change has increased the value of speed in cooperative standards development. The empirical results presented here suggest that a desire for increased 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. This tension between speed and control may explain the increasing popularity of quasi-open standards consortia, such as the Bluetooth Forum, which are typically organized by a handful of large firms.

Finally, the finding that rent-seeking is most pronounced in the upper layers of the IETF protocol stack has implications for the literature on open-ness and innovation. Acemoglu et al. (2010) develop an endogenous growth model where standardization threatens innovation rents, but encourages growth by allowing low-skilled workers to enter the production of standardized goods. Consistent with their central assumption, I find that the standard setting process is more contentious in areas of rapid innovation (i.e. at the top of the stack). However, the relative lack of conflict at lower layers in the IETF protocol stack may be equally significant. In particular, incentives to create new technology can increase with access to a body of knowledge that is codified in well-established standards, just as other "open" institutions have been shown to promote cumulative progress in other domains (e.g. Furman and Stern, 2006; Murray and Stern, 2007). In this view, the Internet's success as a platform for distributed innovation partly reflects the development of many key protocols prior to the era of rapid commercialization.

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. Unlike previous models of coordination delay within SSOs, this theory allows players to make concessions, and suggests that some delay will be efficient when the underlying technology improves over time. The model predicts 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 maneu-

[40] Previous drafts of this paper analyzed the link between quality and delays by examining citations to IETF standards. I find a strong positive correlation between publication delay and citation rates for standards, but not for nonstandards-track RFCs. While this is broadly consistent with the model's prediction that distributional conflict leads to both longer delays and a higher quality cut-off, it could also reflect unobserved variation in the importance of different standards, which leads to both more cites and longer delays.

vering within standard setting committees. In particular, committees with more commercial participants take longer to reach a consensus. To control for omitted variables, I focus on the difference in time-to-consensus for standards versus a nonstandards-track control sample, and use matching methods, instrumental variables and a switching model to guard against bias from endogenous selection onto a particular track. Overall, the empirical results show that distributional conflicts caused by rapid Internet commercialization led to a slower standard setting process. The results also suggest that rent-seeking has a larger impact on coordination delays in more recent years, and in technology areas where innovation is more rapid.

These findings have several implications. First, SSOs reach consensus more quickly in an environment that is free of the pressures created by imminent commercialization. While this is not a complete answer to the larger normative question of when committees will outperform markets as an institution for governing shared technology platforms, it does highlight coordination delays as one important cost of using the consensus process. The results 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 more disputes over compatibility standards and individual firms' rights in the underlying technology.

More broadly, this paper highlights the role of rent-seeking as a source of coordination costs. While established standards promote inter-operability and a division of innovative labor, the process of creating a new component or interface is likely to be contested when it alters the distribution of rents among industry participants. This insight can be applied outside the realm of compatibility standards to a host of coordination problems, such as efforts to create certification programs, codes of conduct, or rules for governing shared natural resources. However, while this paper highlights the link between distributional conflict and bargaining delays, the relevant margin in other settings might be *ex post* compliance, forum shopping or the proliferation of competing standards.

References

Abbate, J. (1999). *Inventing the Internet.* Cambridge, Mass: MIT Press.

Acemoglu, D., G. Gancia, and F. Zilibotti (2010). Competing Engines: of Growth: Innovation and Standardization. *NBER Working paper 15958*.

Angrist, J. D. and S. J. Pischke (2009). *Mostly Harmless Ecnometrics: An Empiricist's Companion.* Princeton University Press.

ANSI (2006). ANSI essential requirements: Due process requirements for American National Standards. http://www.ansi.org (accessed January 2009).

Becker, G. S. and K. M. Murphy (1992). The division-of-labor, coordination costs and knowledge. *Quarterly Journal of Economics 107*(4), 1137–1160.

Bekkers, R., G. Duysters, and B. Verspagen (2002). Intellectual property rights, strategic technology agreements and market structure: The case of GSM. *Research Policy 31*, 1141–1161.

Berners-Lee, T. and M. Fischetti (1999). *Weaving the Web : The original design and ultimate destiny of the World Wide Web by its inventor* (1st ed.). San Francisco: HarperSanFrancisco.

Besen, S. M. and J. Farrell (1994). Choosing how to compete - strategies and tactics in standardization. *Journal of Economic Perspectives 8*(2), 117–131.

Besen, S. M. and L. Johnson (1988). Compatibility standards, competition and innovation in the broadcasting industry. *RAND Study R-3453-NSF. (http://www.rand.org/pubs/reports/R3453.1).*

Besen, S. M. and G. Saloner (1989). The economics of telecommunications standards. In R. Crandall and K. Flamm (Eds.), *Changing the Rules: Technological Change, International Competition, and Regulation in Telecommunications*, pp. 177–220. Washington: Brookings.

Boudreau, K. (2010). Open Platform Strategies and Innovation: Granting Access versus Devolving Control. *Management Science forthcoming.*

Bradner, S. (1996, October). Rfc 2026: The internet standards process – revision 3. Technical report, Internet Engineering Task Force.

Bradner, S. (1997). Key words for use in RFCs to indicate requirement levels. Rfc 2029, Internet Engineering Task Force.

Bradner, S. (1998). IETF Working Group Guidelines and Procedures. Rfc 2418, Internet Engineering Task Force.

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

Brim, S. (2004, February). Guidelines for working groups on intellectual property issues (rfc 3669). Technical report, Inetrent Engineering Task Force.

Brown, S. A. (1997). *Revolution at the Checkout Counter: The Explosion of the Bar Code*. Harvard University Wertheim Committee.

Busch, L.-A. and Q. Wen (1995). Perfect equilibria in a negotiation model. *Econometrica 63*(3), 545–65.

Cargill, C. (2001). Evolutionary pressures in standardization: Considerations on ANSI's national standards strategy. *Testimony Before the U.S. House of Representatives Science Committee. (http://www.sun.com/software/standards/Testimony-Final-reformatted.pdf)*.

Chiao, B., J. Lerner, and J. Tirole (2007). The rules of standard setting organizations: An empirical analysis. *RAND Journal of Economics 38*(4), 905–930.

David, P. A. and S. Greenstein (1990). The economics of compatibility standards: An introduction to recent research. *Economics of Innovation and New Technology 1*, 3–42.

Davies, E. (2004). IETF problem statement. Request for Comments 3774, Internet Engineering Task Force, http://www.ietf.org/rfc/rfc3774.

Farrell, J., J. Hayes, C. Shapiro, and T. Sullivan (2007). Standard setting, patents and hold-up. *Antitrust Law Journal 74*, 603–670.

Farrell, J. and G. Saloner (1986). Installed base and compatibility - innovation, product preannouncements, and predation. *American Economic Review 76*(5), 940–955. Times Cited: 200 Article English Cited References Count: 11 F2246.

Farrell, J. and G. Saloner (1988). Coordination through committees and markets. *RAND Journal of Economics 19*(2), 235–252.

Farrell, J. and C. Shapiro (1992). Standard setting in high-definition television. *Brookings Papers on Economic Activity*, 1–93.

Farrell, J. and T. Simcoe (2008). Choosing the rules for formal standardization. *Unpublished manuscript*.

Foray, D. (1995). Coalitions and committees: How users get involved in information technology standardization. In R. Hawkins, R. Mansell, and J. Skea (Eds.), *Standards, Innovation and Competitiveness*, pp. 192–212. Edward Elgar.

Furman, J. and S. Stern (2006). Climbing atop the shoulders of giants: The impact of institutions on cumulative research. *NBER Working Paper No. 12523*.

Heckman, J. (1976). The Common Structure of Statistical Models of Truncation, Sample Selection, and Limited Dependent Variables, and a Simple Estimator for such Models. *Annals of Economic and Social Measurement 5*, 475–492.

IETF (1990, June). Internet monthly report. Technical report, Internet Engineering Task Force.

Jones, B. F. (2008). The Knowledge Trap: Human Capital and Development Reconsidered. *NBER Working Paper #14138*.

Katz, M. L. and C. Shapiro (1985). Network externalities, competition and compatibility. *American Economic Review 75*, 424–440.

Kennan, J. and R. Wilson (1993, March). Bargaining with private information. *Journal of Economic Literature 31*(1), 45–104.

Lehr, W. (1995). Compatibility standards and interoperability: Lessons from the internet. In B. Kahin and J. Abbate (Eds.), *Standards Policy for Information Infrastructure*. Cambridge, MA: MIT Press.

Leiponen, A. (2008). Competing through cooperation: Standard setting in wireless telecommunications. *Management Science 54*(2), 1904–1919.

Lemley, M. (2002). Intellectual property rights and standard setting organizations. *California Law Review 90*, 1889–1981.

Lemley, M. (2007). Ten things to do about patent holdup of standards (and one not to). *Boston College Law Review 48*, 149–168.

Lerner, J. and J. Tirole (2006). A model of forum shopping. *American Economic Review 96*(4), 1091–1113.

Levinson, M. (2006). *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger*. Princeton University Press.

Lippman, S. A. and J. J. McCall (1976). The economics of job search: A survey. *Economic Inquiry 4*, 17–37.

Merlo, A. and C. Wilson (1995). A stochastic model of sequential bargaining with complete information. *Econometrica 63*(2), 371–399.

Murray, F. and S. Stern (2007). Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis. *Journal of Economic Behavior and Organization 63*(4), 648–687.

Myatt, D. 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.

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action* . Cambridge University Press.

Postel, J. (1995). Internet official protocol standards (RFC 1800). Technical report, Internet Engineering Task Force.

Rosenbaum, P. R. and D. B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika 70*, 41–55.

Rosenkopf, L., A. Metiu, and V. P. George (2001). From the bottom up? technical committee activity and alliance formation. *Administrative Science Quarterly 46*, 748–772.

Russell, A. L. (2006). 'Rough Consensus and Running Code and the Internet-OSI Standards War. *IEEE Annals of the History of Computing July-September*, 48–61.

Rysman, M. and T. Simcoe (2008). Patents and the performance of voluntary standard setting organizations. *Management Science 54*(11), 1920–1934.

Saltzer, J. H., D. P. Reed, and D. D. Clark (1984). End-to-end arguments in system design. *ACM Transactions on Computer Systems 2*(4), 277–288.

Shapiro, C. and H. R. Varian (1998). *Information Rules : A Strategic Guide to the Network Economy*. Boston, Mass.: Harvard Business School Press.

Stock, J. and M. Yogo (2005). Testing for Weak Instruments in Linear IV Regression. *NBER Technical Working Paper 284*.

Updegrove, A. (2007, March). Intellectual property rights and standard setting. *Consortium Stadnards Bulletin 6*(3).

Vytlacil, E. (2002). Independence, Monotonicity, and Latent Index Models: An Equivalence Result. *Econometrica 70*(1), 331–341.

Waguespack, D. and L. Fleming (2009). Scanning the commons? evidence on the benefits to statrups of participation in open standards development. *Management Science 55*, 210–223.

Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. MIT Press.

Yildiz, M. (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.
	Outcomes & Instruments			
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
log(ID Mail)	Log count of emails mentioning ID	ID	2.98	1.25
S-Track	Dummy for Proposed Standard	ID	0.40	0.49
N-Track	Dummy for Nonstandards-track RFC	ID	0.21	0.41
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
Keywords	Dummy for S-Track Keywords > 0	ID	0.68	0.47
log(KeyCount)	Log count of S-Track keywords	ID	1.74	1.54
	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
Org-edu	Maximum of Dot-org and Dot-edu	ID	0.25	0.43
	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
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
Dot-gov	Dummy for author from dot-gov TLD	ID	0.05	0.21
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

	Full Sample			Matched Sample†		
Variable	Nonstds Track	Proposed Standards	P-value (ST=NST)	Nonstds Track	Proposed Standards	P-value (ST=NST)
Total Days	568.16	702.21	0.00	545.69	707.13	0.00
Versions	5.73	7.60	0.00	5.50	7.68	0.00
log(ID Mail)	3.15	3.51	0.00	3.16	3.51	0.00
RFC Cites	5.02	10.88	0.00	5.37	10.02	0.00
Patent Cites	1.58	3.20	0.00	1.80	3.06	0.05
Article Cites	1.62	2.09	0.41	1.60	2.00	0.59
Keywords	0.59	0.72	0.00	0.60	0.71	0.00
log(KeyCount)	1.41	1.92	0.00	1.42	1.84	0.00
Suit-share	71.79	73.59	0.08	73.47	73.29	0.87
log(Patents)	7.39	7.51	0.53	7.21	7.46	0.27
Dot-org	0.08	0.08	0.94	0.09	0.07	0.43
Dot-edu	0.23	0.17	0.03	0.18	0.17	0.79
Org-edu	0.28	0.24	0.16	0.24	0.24	0.85
Draft Cohort	1998.90	1998.76	0.37	1998.70	1998.71	0.98
log(WG Mail)	5.44	5.85	0.00	5.63	5.83	0.10
log(Cum Mail)	6.29	6.84	0.00	6.54	6.85	0.01
log(Drafts)	1.79	1.94	0.00	1.85	1.92	0.09
log(Cum Drafts)	2.36	2.66	0.00	2.48	2.64	0.01
log(Members)	2.13	2.22	0.07	2.13	2.19	0.29
log(Filesize)	10.21	10.31	0.09	10.18	10.26	0.23
ID Sponsors	2.13	2.00	0.13	2.02	2.02	0.98
WG Chair	0.42	0.46	0.19	0.41	0.45	0.31
Dot-gov	0.06	0.06	0.85	0.06	0.06	0.68
Dot-com	0.85	0.88	0.23	0.86	0.87	0.54
Hotelling's T^2 (P-value)		0.00			0.23	
Obs. (Standards)		773			662	
Obs. (Nonstds)		393			290	

Notes: Both samples exclude all RFCs with Total Days > 2007 (5.5 years) to correct for right-truncation of DV in OLS models. †Section 4.3 and Appendix provide details on propensity-score matching.

Table 4: Conflict, Concessions and Coordination Delay

	Dependent Variable = Total Days				
	(1)	(2)	(3)	(4)	(5)
Sample & Specification	Full OLS	Matched OLS	Matched OLS	Full IV†	S-Track RE‡
Suit-share * S-track	7.8 (1.9)***	6.8 (1.9)***	5.4 (2.1)*	12.3 (5.9)*	4.2 (1.5)**
log(Patents) * S-track	6.5 (9.6)	0.6 (10.6)	-1.1 (11.3)		18.6 (6.6)**
Org-edu * S-Track	155.7 (65.7)*	153.6 (73.2)*	92.2 (77.8)	263.7 (171.2)	86.7 (32.6)**
Suit-Share	-3.3 (1.6)*	-1.3 (1.5)	0.7 (2.7)	-6.4 (4.4)	
log(Patents)	4.9 (7.5)	3.6 (8.1)	10.9 (12.0)	11.2 (6.1)+	
Org-edu	-75.0 (59.5)	-60.3 (69.6)	-46.8 (81.2)	-144.8 (119.7)	
	Additional Controls & Statistics				
Additional Controls*	10 [0.00]	10 [0.00]	10 [0.10]	5 [0.00]	5 [0.00]
Draft Cohort FEs	9 [0.00]	9 [0.00]	9 [0.00]	9 [0.00]	9 [0.00]
Technology Class FEs	5 [0.00]	5 [0.01]		5 [0.00]	5 [0.01]
Working Group FEs			151 [0.00]		
Hausman Test			28.42 [0.34]		13.95 [0.67]
Endogeneity Test				1.70 [0.64]	
A-P F-Stat (Suit-share)				43.46	
A-P F-Stat (Org-edu)				50.23	
A-P F-Stat (S-track)				11.53	
Obs. (Standards)	773	662	662	773	773
Obs. (Nonstandards)	393	290	290	393	0
Model dof	31	31	25	26	22
R-Squared	0.18	0.20	0.09	0.07	

+10% *5% **1% ***0.1% significance levels (clustered on Working Group).

Notes: All models contain an S-Track indicator. †Model (4) is a GMM-IV regression that instruments for S-Track and all interactions with S-track. The instruments are fitted values from a probit selection equation that includes Keywords and log(KeyCount), interacted with all exogenous variables (see text for details). ‡Model (5) is a random effects specification. Hausman test null hypothesis is that random effects estimates are consistent (p-value in brackets). *For controls and fixed effects, the table indicates number of parameters and a p-value for joint significance. Additional controls are main effects and S-Track interactions for: log(Filesize), log(Drafts), log(Cum Mail), 1[ID Sponsors=2], and 1[ID Sponsors>2]. See the text for additional details.

Table 5: Conflict and Delay: Switching Model

Model	Switching Model Full Sample (RFCs) (1)			Switching Model All Internet Drafts (2)		
Dependent Variable	S-Track log(Days)	N-Track log(Days)	Selection	S-Track log(Days)	N-Track log(Days)	Selection
log(Suit-share)	0.37	-0.39	0.21	0.27	-0.52	0.22
	(0.14)***	(0.17)*	(0.27)	(0.17)	(0.18)***	(0.25)
log(Patents)	0.02	0.02	-0.03	0.01	0.01	-0.03
	(0.01)	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)
Org-edu	0.09	-0.11	-0.03	0.05	-0.13	-0.05
	(0.06)	(0.11)	(0.11)	(0.06)	(0.11)	(0.11)
Keywords			0.19			0.22
			(0.14)			(0.15)
log(KeyCount)			0.08			0.08
			(0.05)			$(0.05)^{+}$
Covariance ($\sigma_{\nu x}/\sigma^2_{\varepsilon^x}$)	0.14	-0.26		0.13	0.07	
	(0.14)	(0.19)		(0.16)	(0.30)	
	Additional Controls & Statistics					
Additional Controls†	15 [0.00]			15 [0.00]		
Technology Class FEs	10 [0.04]			15 [0.00]		
Draft Cohort FEs	18 [0.00]			30 [0.00]		
Obs. (Standards)	773			935		
Obs. (Nonstandards)	393			474		
Obs. (Expired)	0			820		
Obs. (Censored)	0			58		
	Joint Hypothesis Tests‡					
Suit-share Equality	$\chi^2_1 = 11.72$ [0.00]			$\chi^2_1 = 9.21$ [0.00]		
Org-edu Equality	$\chi^2_1 = 2.83$ [0.09]			$\chi^2_1 = 2.31$ [0.13]		
Weak IVs (Keywords = 0)	$\chi^2_2 = 13.16$ [0.00]			$\chi^2_2 = 20.67$ [0.00]		
Exog. Selection ($\sigma_{\nu x} = 0$)	$\chi^2_2 = 2.18$ [0.34]			$\chi^2_2 = 0.63$ [0.73]		

$^{+}$10% *5% **1% ***0.1% significance levels (clustered on Working Group).

Notes: Model (1) is estimated on the Full Sample of RFCs described in the left panel of Table 3. Model (2) is estimated on all IDs submitted to the IETF between 1993 and 2003. Model (1) constrains Technology Class and Draft Cohort effects to be equal for standards and nonstandards, while model (2) does not. †For controls and fixed effects, the table indicates number of parameters and a p-value for joint significance. Additional controls are main effects and S-Track interactions for: log(Filesize), log(Drafts), log(Cum Mail), 1[ID Sponsors=2], and 1[ID Sponsors>2]. ‡For joint hypothesis tests, the number in brackets is a p-value. See Section 4.1 for discussion of the switching model.

Table 6: Conflict and Delay: OLS Interaction Effects

	Dependent Variable = Total Days					
Model	(1)	(2)	(3)	(4)	(5)	(6)
Dot-org†	9.4 (3.7)*	10.5 (6.1)+				
Dot-edu†	-0.1 (2.4)	1.1 (3.8)				
Applications Area†			9.0 (2.3)***	7.5 (3.4)*		
Transport Area†			5.0 (2.7)+	11.7 (3.2)***		
Internet Area†			3.5 (3.2)	1.6 (4.3)		
Routing Area†			-1.1 (3.3)	6.0 (5.6)		
Security Area†			-0.8 (5.9)	6.5 (6.6)		
Operations Area†			3.2 (3.5)	7.9 (6.0)		
Cohort [93,94]†					0.1 (3.7)	-1.6 (4.9)
Cohort [95,96]†					4.8 (2.5)+	5.6 (4.5)
Cohort [97,98]†					0.9 (2.6)	5.7 (4.1)
Cohort [99,00]†					6.4 (2.4)**	7.5 (3.4)*
Cohort [00,01]†					5.7 (3.5)	12.0 (3.4)***
	Additional Controls & Statistics					
Suit-share * Dot-Org/Edu‡		2 [0.87]				
Suit-share * Technology FEs‡				6 [0.05]		
Suit-share * Cohort FEs‡						5 [0.25]
Obs. (Standards)	773	773	773	773	773	773
Obs. (Nonstandards)	0	393	0	393	0	393
Model d.o.f.	24	49	26	53	25	46
R-Squared	0.22	0.20	0.22	0.21	0.21	0.19

$^{+}$10% *5% **1% ***0.1% significance levels (clustered on Working Group).

Notes: †Coefficients from models (1), (3) and (5) are two-way interactions with Suit-share. Coefficients from models (2), (4) and (6) are three-way interactions with Suit-share and S-Track. All models include Technology Class and Publication Cohort fixed effects, along with main effects and and a full set of two-way interactions. Additional controls are main effects and S-Track interactions for: log(Filesize), log(Drafts), log(Cum Mail), 1[ID Sponsors=2], and 1[ID Sponsors>2]. ‡Number of parameters and p-value for joint significance of nonstandards-track (two-way) interaction effects. See text for additional details.

Appendix: Proofs

It is useful to have an explicit solution for concessions $b(x)$. Since constraints (2) and (3) bind at $x = q^*$, one can solve for βV and substitute back into $b(x) = \max\{0, \frac{\beta V - \pi(x,0)}{\gamma}\}$ to show that $b(x) = \max\left\{0, \frac{\gamma\pi(q^*,1)+k\pi(q^*,0)-(k+\gamma)\pi(x,0)}{\gamma(k+\gamma)}\right\}$.

Section 2.2: Equilibrium Results

Result 1: *There is a unique symmetric equilibrium.*

Proof: Substitute for $b(x)$ in (4). The left side of this equation is strictly decreasing in q^* and approaches 0 as $q^* \to \infty$. The right side of (4) is bounded below by $(k+1)c > 0$ and is increasing in q^*, so the two sides cross exactly once between 0 and ∞. $\square$

Result 2: *Delays are inefficient (q^* is above the first best) if and only if $\omega > 0$ and $\gamma < 1$.*

Proof: When $\omega = 0$, then $\pi(x, 1) = \pi(x, 0)$ by assumption, so $b(x) = 0$ for all x, and terms containing $b(x)$ drop out of equation (4). When $\gamma = 1$ (transferable utility), then $b(x) > 0$ but the same terms drop out of (4). Without the terms containing $b(x)$, equation (4) is a joint surplus maximizing (i.e. socially efficient) cutoff rule, so both conditions are necessary.

To see that $\omega > 0$ and $\gamma < 1$ are sufficient conditions, note that $\omega > 0$ implies $b(q^*) > 0$. When $b(q^*) > 0$ and $\gamma < 1$ the left (right) side of equation (4) is strictly larger (smaller) than if $\omega = 0$ or $\gamma = 1$. Since the left (right) side of (4) is decreasing (increasing) in q^*, the equilibrium cut-off must be above the first best whenever both inequalities are satisfied. $\square$

Section 2.3: Comparative Statics

Prediction 1: Expected coordination delays are (weakly) increasing in ω if opportunity costs also decline ($S_\omega(x) \leq 0$); or players are sufficiently patient ($\beta = 1$); or there is no redistribution ($\pi_\omega(x, 0) = 0$) and concessions are sufficiently costly ($\gamma = 0$)

Proof: Define $\widetilde{q}$ as the worst proposal accepted without concessions, so $\pi(\widetilde{q}, 0; \omega) = \beta V$. Substitute for $b(x)$ in (4), and totally differentiate that expression with respect to ω. Using the assumption that marginal returns to quality $\pi_q(x, a; \omega)$ are constant, we arrive at

$$\frac{dq^*}{d\omega} S_q(q^*)[\beta^{-1} - G(q^*)] = (1 - \beta^{-1})S_\omega(q^*) + k(\gamma - 1)[G(\widetilde{q}) - \beta^{-1}]b_\omega(q^*)$$

Since $S_q(q^*)[\beta^{-1} - G(q^*)] > 0$, we have $\frac{dq^*}{d\omega} > 0$ as long as the right side of this expression is positive. Note that $b_\omega(q^*) > 0$, since π has increasing differences in (ω, a) by assumption. Thus, $\frac{dq^*}{d\omega} > 0$ whenever $S_\omega(x) \leq 0$ or $\beta = 1$. Finally, when $\pi_\omega(q^*, 0) = 0$, the right side of the expression approaches $\pi_\omega(q^*, 1)[1 - G(q^*)] > 0$ as $\gamma \to 0$. $\square$

Prediction 2: Expected coordination delays are (weakly) decreasing in γ.

Proof: Totally differentiating (4) with respect to γ yields:

$$\frac{dq^*}{d\gamma}S_q(q^*)[\beta^{-1} - G(q^*)] = \frac{k(k+1)}{k+\gamma}[G(\widetilde{q}) - \beta^{-1}]b(q^*) < 0.$$

Prediction 3: The positive impact of increased distributional conflict on coordination delays is decreasing in γ: $\frac{d^2q^*}{d\gamma d\omega} < 0$.

Proof: First note that $\frac{d\widetilde{q}}{d\gamma} > 0$. This can be shown by solving for βV and totally differentiating the identity $\pi(\widetilde{q}, 0; \omega) = \beta V$, or by noting that V increases with γ (by the envelope theorem) and π is increasing in q. We also know from the proof of Prediction 1 that

$$\frac{dq^*}{d\omega}S_q(q^*)[\beta^{-1} - G(q^*)] = (1-\beta^{-1})S_\omega(q^*) + k(\gamma - 1)[G(\widetilde{q}) - \beta^{-1}]b_\omega(q^*)$$

Totally differentiating this equation with respect to γ, using the assumption that the marginal returns to quality are constant (which implies that $\frac{\partial^2 S(x)}{\partial q^2} = \frac{\partial^2 b(x)}{\partial q^2} = 0$), and collecting terms leads to the following equation

$$\frac{d^2q^*}{d\gamma d\omega}S_q(q^*)[\beta^{-1} - G(q^*)] = g(q^*)\frac{dq^*}{d\gamma}\frac{dq^*}{d\omega}S_q(q^*) + k[G(\widetilde{q}) - \beta^{-1}]b_\omega(q^*)\left[1 + \frac{1-\gamma}{(k+\gamma)^2}\right] + k(\gamma - 1)g(\widetilde{q})\frac{d\widetilde{q}}{d\gamma}b_\omega(q^*)$$

Since $b_\omega(q^*) > 0$, $S_q(x) > 0$, $\frac{dq^*}{d\gamma} > 0$, $\frac{dq^*}{d\omega} < 0$ and $\frac{d\widetilde{q}}{d\gamma} > 0$, each term on the right hand side of the equation is negative. This implies that $\frac{d^2q^*}{d\gamma d\omega} < 0$. $\square$

DATA APPENDENDICES (Not for Publication)

A1 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. These searches were conducted in July 2008.

Independent Variables

Internet Drafts

The IETF's file naming conventions were used to match each Internet Draft to a particular Working Group. In a few cases where individual ID's were later adopted by a Working Group, the two series were matched by hand. 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. (Mark Overmeer's Mail::Box perl modules were invaluable for processing the email data.) 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 organizational affiliation of each Internet Draft author is identified using an email address from inside the ID. 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.[41]

I successfully matched 193 organizations to an USPTO 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.

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

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.

A2 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 A1. Figure A1 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 952 out of the original 1,166 RFCs, or 82 percent of the initial sample.

Figure A1: Propensity Score Distribution

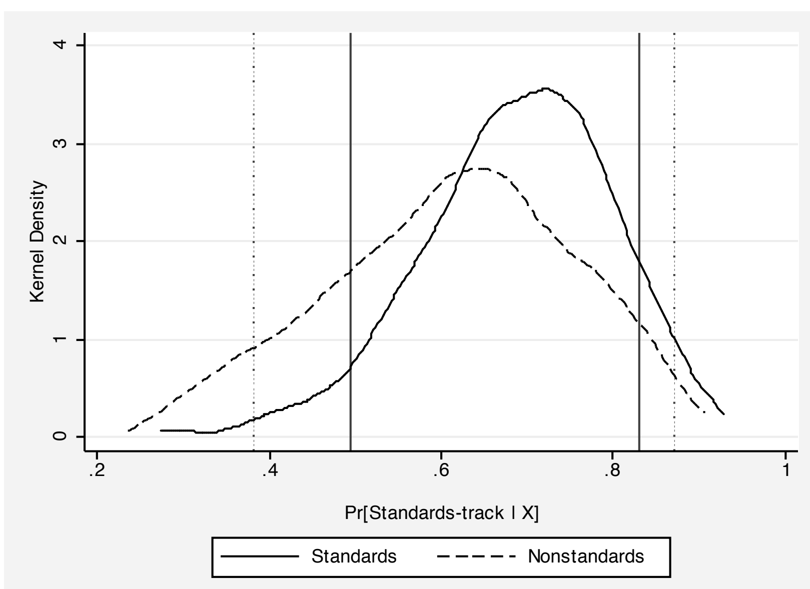

Table A1: Propensity Score Probit

	Probit Regression DV = S-Track	
	Coefficients	Marginal Effects
Suit-share	0.01 (0.00)**	0.00 (0.00)**
Lag Suit-share	-0.01 (0.00)**	-0.00 (0.00)**
log(Patents)	-0.04 (0.02)*	-0.02 (0.01)*
log(Drafts)	-0.90 (0.45)*	-0.32 (0.16)*
log(Cum Drafts)	0.65 (0.29)*	0.23 (0.10)*
log(Members)	0.11 (0.08)	0.04 (0.03)
log(Email)	-0.01 (0.05)	-0.00 (0.02)
log(Cum Email)	0.06 (0.12)	0.02 (0.04)
log(Cum Drafts)^2	-0.10 (0.05)*	-0.04 (0.02)*
log(Cum Email)^2	-0.00 (0.01)	-0.00 (0.00)
log(Drafts)^2	0.23 (0.11)*	0.08 (0.04)*
Sponsors	-0.06 (0.03)+	-0.02 (0.01)+
WG Chair	0.19 (0.09)*	0.07 (0.03)*
log(Filesize)	0.14 (0.05)**	0.05 (0.02)**
Dot-org	0.17 (0.16)	0.06 (0.05)
Dot-edu	-0.17 (0.10)+	-0.06 (0.04)
Dot-gov	-0.09 (0.17)	-0.03 (0.06)
Tech Class Effects	5 [0.01]	
ID Cohort Effects	2 [0.34]	
WG Cohort Effects	2 [0.22]	
Pr(S-Track=1)	0.67	
Observations	1166	
Model dof	26	

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

A3 Additional Descriptive Statistics

Table A2: 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.

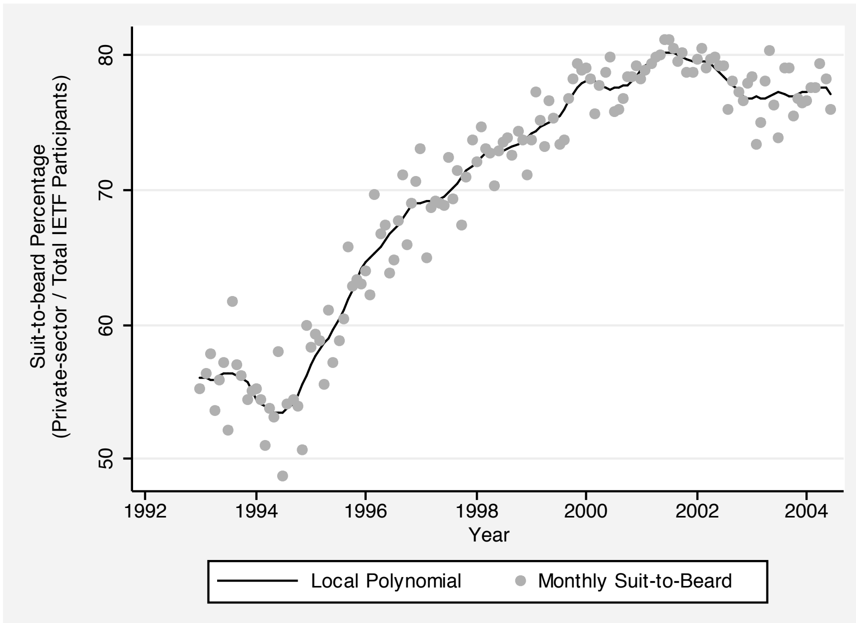

Figure A2: IETF Commercialization (Monthly Suit-share)

Table A3: 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 A4: 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.

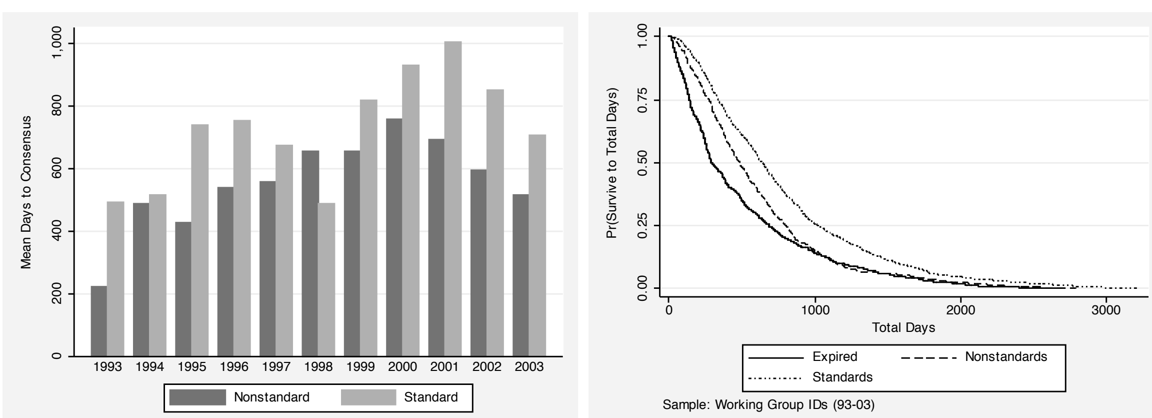

Figure A3: Mean Duration by Cohort (left) & Survival Curves (right)

A4 Switching Model

I estimate a switching model with partially observed regimes, which is based on the endogenous selection model introduced by Heckman (cites). The model has three equations:

$$T_i^m = X_i\beta^m + \varepsilon_i^m$$
$$S_i = 1[Z_i + \nu_i]$$

where $m \in \{s, n\}$ indexes standards and nonstandards respectively, and the unobservables have a trivariate normal distribution $(\varepsilon^s, \varepsilon^n, \nu) \sim N(0, \Sigma)$. 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

$$\eta_{im} = \frac{\gamma Z_i + \rho_{mv}\varepsilon_i^m/\sigma_m}{\sqrt{1-\rho_{mv}^2}}$$

where $\rho_{mv} = \sigma_{mv}/\sigma_m$ is the coefficient of correlation between ε^m and ν. Since the unobservables 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 $\sum_i \ln(Pr(S_i, T_i))$. Code for estimating this model in Stata was adapted from the movestay routine developed by Lokshin and Sajaia (2004) and is available on the author's web site.

A5 Additional Results

Table A5: 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	5.7	9.1	15.9	4.8	3.7	10.5	11.5
	(2.2)**	(2.9)**	(11.2)	(1.7)**	(4.3)	(5.7)+	(8.8)
Suit-share		-3.3		-1.8	-0.1		0.5
		(2.4)		(2.8)	(4.0)		(7.3)
	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				
F-Test of Excluded IVs							
Suit-share * S-track	203.5 [0.00]	286.5 [0.00]	2.75 [0.07]				
Suit-share		147.5 [0.00]					
	Additional Controls & 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)	773	667	737	385	385	118	118
Obs (N-track)	0	289	0	281	281	0	51

+10% *5% **1% ***0.1% significance levels (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 main effects and S-Track interactions for: log(Filesize), log(Drafts), log(WG Mail), log(Cum Mail), Dot-gov, 1[ID Sponsors=2], and 1[ID Sponsors>2]. Additional controls in (4) and (5) are main effects and S-Track interactions for: log(Filesize), log(Drafts), log(WG Mail), log(Cum Mail), Dot-gov, 1[ID Sponsors=2], and 1[ID Sponsors>2].

Table A6: Diff-in-diffs: Alternative Measures and Specifications

Model	Big-Firm Effects		ID Suit-Share		Poisson QML		Cox Hazard		75th Quantile Regression	
Suit-share * S-track	7.34 (2.00)***	6.09 (2.18)**			0.66 (0.19)***	0.51 (0.22)*	-1.79 (0.47)***	-2.10 (0.45)***	6.58 (2.27)**	
log(Patents) * S-track	4.42 (10.50)	-1.59 (11.65)	2.03 (10.56)	3.88 (11.13)	0.00 (0.02)	0.00 (0.02)	0.03 (0.02)	-0.00 (0.03)	9.16 (11.18)	
Org-edu * S-Track	171.43 (75.87)*	76.16 (84.10)	117.54 (73.08)	85.94 (74.13)	0.25 (0.13)+	0.16 (0.14)	-0.33 (0.16)*	-0.13 (0.17)	87.70 (70.01)	
Suit-share	-1.46 (1.69)	0.13 (2.77)			-0.16 (0.16)	0.06 (0.25)	0.98 (0.38)*	0.17 (0.54)		-2.21 (3.33)
log(Patents)	3.55 (8.46)	11.71 (12.53)	0.72 (8.05)	13.22 (12.35)	0.01 (0.01)	0.01 (0.02)	-0.04 (0.02)*	-0.03 (0.03)		-13.75 (19.41)
Org-edu	-45.58 (80.12)	-48.14 (95.19)	-18.96 (66.93)	-4.64 (76.06)	-0.11 (0.13)	-0.09 (0.15)	0.19 (0.14)	0.06 (0.15)		-57.64 (99.65)
ID Suit-share * S-track			3.18 (1.41)*	3.09 (1.34)*						
ID Suit-share			1.84 (1.15)	3.48 (1.42)*						
Tech Class Effects	Y	na	Y	na	Y	na	Y	na	Y	Y
WG Effects	N	Y	N	Y	N	Y	N	Y	N	N
PubCohort Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Large Firm Effects	Y	Y	N	N	N	N	N	N	N	N
Additional Controls†	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Obs. (Standards)	662	662	662	662	662	662	935	935	773	0
Obs. (Nonstds)	290	290	290	290	290	290	474	474	0	393
Obs. (Censored)							878	878		

$^{+}$10% *5% **1% ***0.1% significance levels (clustered on Working Group).

Notes: Big-Firm, ID Suit-share and Poisson models are estimated on the Matched Sample. Cox Hazard model assumes all expired proposals are censored standards-track observations. †Additional controls in all models are S-track are main effects and S-Track interactions for: log(Filesize), log(Drafts), log(Cum Mail), 1[ID Sponsors=2], and 1[ID Sponsors>2].

Table A7: Publication Delay and Forward Citations

	Poisson Regressions DV = Total Cites					
	(1)	(2)	(3)	(4)	(5)	(6)
log(Suit-share)	-0.11 (0.44)		-0.11 (0.50)	1.48 (0.53)**		1.50 (0.54)**
log(Suit-share) * S-track	0.07 (0.45)		0.07 (0.50)	-0.88 (0.49)		-0.96 (0.51)
log(Total Days)		-0.10 (0.10)	-0.10 (0.09)		-0.06 (0.11)	-0.05 (0.11)
log(Total Days) * S-track		0.36 (0.12)**	0.36 (0.12)**		0.36 (0.12)**	0.35 (0.12)**
log(ID Mail)	0.35 (0.04)***	0.34 (0.05)***	0.34 (0.04)***	0.36 (0.05)***	0.35 (0.05)***	0.34 (0.05)***
log(Cites/Month)$_{-i}$	0.64 (0.07)***	0.64 (0.07)***	0.64 (0.07)***			
	Control Variables & Regression Statistics					
S-track x Cohort FE's	20 [0.00]	20 [0.00]	20 [0.00]	20 [0.00]	20 [0.00]	20 [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	6 [0.00]	6 [0.00]	6 [0.00]	6 [0.00]	6 [0.00]	6 [0.00]
Tech-Class FEs	5 [0.00]	5 [0.00]	5 [0.00]			
WG FEs				Y	Y	Y
Obs. (Standards)	743	743	743	734	734	734
Obs. (Nonstandards)	403	403	403	393	393	393
Model d.o.f.	39	39	41	33	33	35
Pseudo-LogL x 10^3	-8.30	-8.21	-8.21	-5.83	-5.82	-5.75

*5% significance; **1% significance; ***0.1% 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.