Learning When to Quit*

Bernhard Ganglmair† — *University of Texas at Dallas*

Timothy Simcoe‡ — *Boston University & NBER*

Emanuele Tarantino§ — *University of Mannheim & CEPR*

October 5, 2017

Abstract

We propose a model of research and development as a process of experimentation in which researchers repeatedly revise specifications of a project and update their beliefs about the project's type. Only a good project whose type is learned by researchers can generate value. Researchers abandon a project when the opportunity costs of continuing exceed the expected benefits. We estimate the structural parameters of this dynamic optimization problem using a novel data set with information on both successful and abandoned projects from the Internet Engineering Task Force (IETF), an organization that creates and maintains standards necessary for the functioning of the internet. The structural approach allows us to recover researchers' unobserved beliefs and opportunity costs, and answer questions about whether specific rules and institutions encourage "efficient abandonment" of ongoing projects. We find that opportunity costs are decreasing over time, and feedback and comments from the IETF community at large increase the speed at which developers learn whether a project is worth pursuing.

*We thank Marc Rysman, Florian Schütt, Carlos Serrano, and participants at the TILEC conference "Competition, Standardization, and Innovation" (Amsterdam) and the MaCCI conference on Innovation (Bad Homburg) for comments and suggestions, and Ajándék Peák, Sucheth Prasad, and Jacob Walsh for excellent research assistance. Bernhard Ganglmair would like to thank the economics department at Boston University for their hospitality.

†The University of Texas at Dallas, Naveen Jindal School of Management, 800 W. Campbell Rd (SM31), Richardson, TX 75080, USA; Email: ganglmair@utdallas.edu.

‡Boston University, Questrom School of Business, 595 Commonwealth Ave, Boston, MA 02215, USA; Email: tsimcoe@bu.edu.

§University of Mannheim, Department of Economics, L7, 3–5, 68131 Mannheim, Germany; Email: tarantino@uni-mannheim.de.

1 Introduction

There is a vast literature on the economics of innovation, primarily focused on the question of how specific institutions influence the "rate and direction" of technological change (Lerner and Stern, 2010). This broad literature has examined how the patent system, universities, government R&D support, and the norms of "open science" all contribute to the production and dissemination of knowledge. While many studies focus on the behavior of individual scientists and engineers, very few (perhaps none) analyze how these individuals allocate the key research input of *time* in the face of substantial uncertainty. That is the question and contribution at the heart of this paper.

Most researchers have decided to abandon an idea or project at some point.[1] This decision reveals that the perceived opportunity costs of continuing down a particular path exceed the expected benefits. Yet, if the project was started, the expected benefits must have exceeded the opportunity costs at the outset. This suggests that researchers learn about the expected costs and benefits of a line of research during a project's development. Because beliefs and opportunity costs are not directly observable, we will need a model in order to answer questions about whether specific rules and institutions encourage researchers' efficient abandonment of the project.

Our proposed model is motivated by several stylized facts about the process of collaborative R&D that we observe at the Internet Engineering Task Force (IETF). The IETF is an organization that develops and maintains the core technological standards for the Internet, and it provides an ideal setting

[1] Henceforth, we use the terms idea, project and proposal interchangeably.

for studying this problem because of its highly transparent processes. Using data on about 16,000 IETF projects initiated between 1996 and 2009, we show that many ideas fail quickly. In particular, the hazard of abandonment drops sharply, and the hazard of publication grows gradually with the number of revisions made to a particular idea. Secondly, we find that increased communication (via email) is associated with faster failure, and slower publication. And third, we observe a strong positive and monotonic relationship between the number of revisions to an idea, and the number of U.S. patent citations that it subsequently receives.

Our model of Bayesian learning combines a one-armed bandit problem with a more traditional optimal stopping problem to capture two different phases of the research process.[2] We assume there are two types of idea, good (publishable) and bad (doomed to fail). In the first phase of the process, a team of researchers runs a sequence of experiments striving to learn whether a project is of the good type. A project's true type is realized only if there is a *breakthrough* leading to a consensus that the project merits publication. In the second phase, conditional on reaching consensus, the team continues developing the project to bring it to completion.[3]

The key parameters in our model are the players' beliefs about the distribu-

[2] Bergemann and Välimäki (2008) provide a survey of the economics literature on bandit problems. For earlier applications of bandit models to economics, see Rothschild (1974)

[3] Our definition of consensus is different from that in the IETF. There, a project is published as RFC (Request for Comments) when a working group chair finds that "rough consensus" has been reached. In our framework, consensus is not on the final version (after the second phase) of the project but on the type of the project (after the first phase). Consensus in our empirical context means that researchers observe the good type of the project (e.g., after receiving a sufficient number of positive signals from the community) and anticipate that it will be published after further revisions during the second phase.

tion of good and bad projects, the rate at which they learn and the opportunity costs of continuing a project. We recover these parameters by maximum likelihood estimation of the learning model using data from the IETF. For each project submitted to the IETF during our sample period, we observe when it was initiated, its outcome (published or abandoned), the number of revisions submitted by the author team, the size of the team, the extent of communication about the project (e-mails/version), and the number of citations of projects in U.S. patents (from 1976 to 2015). We begin the estimation by fitting a model of patent citations conditional on the number of revisions to a successful project. Given this payoff function, and assuming a set of independently distributed cost shocks, it is possible to solve the learning model backwards (recursively) to obtain the likelihood of the data for a given set of parameters. Intuitively, the learning and cost parameters are identified by the rate at which IETF projects are published and abandoned, as well as the overall share of projects that reach each end point.

Our estimates imply that the marginal opportunity costs of an additional revision are decreasing and convex in the version number, with a steep initial decline. The estimated opportunity costs are higher when projects are initially sanctioned by an IETF working group and when there are fewer researchers on the team. This implies that, while we are agnostic about the functioning of collaboration within the team, larger teams face lower costs.

We find that projects initiated by IETF working groups have a higher rate of learning than the average project. Researchers in working groups learn the type of the project faster and abandon bad projects faster than researchers of projects that are initially not sanctioned by a working group. One explanation for this is that projects sanctioned by working groups receive more attention and feedback than outside projects. We also further find that researchers' prior beliefs that a project is good are higher for projects initiated by a working group.

Finally, we find that projects that have triggered more discussion and received more comments by the IETF community (measured in terms of e-mails per version sent in response to a new version) exhibit a higher rate of learning. This suggests that attention and feedback from the IETF community at large and communication with other researchers increase the speed at which researchers learn the type of their project. With a higher rate of learning, research teams abandon projects faster in phase one. This means that communication results in a more efficient process because bad projects are abandoned earlier. But it also means that researchers are more impatient, thus abandoning good projects for which they may otherwise observe a breakthrough.

We use our structural model of learning in collaborative R&D to calculate two counterfactuals. For our first counterfactual, we treat the agents' shared beliefs about the distribution of project types as an institutional variable. The prior probability is the researchers' expectations that a consensus can be reached and the project will eventually be published as an RFC. In our second counterfactual, we consider the effect of imposing a deadline on the publication process. One of the interesting features of our model is that it implies some share of IETF projects are "false negatives" that could achieve consensus, but fail to do so in time, and are abandoned. We focus on both publication timing and the false negative rate in our counterfactual analysis.

Our paper makes a number of contributions. We construct a novel data set with information on both successful (or published) and abandoned R&D projects. This information provides us with a unique opportunity to study (i) the speed of learning in R&D and (ii) how specific rules and institutions influence the speed of learning about bad approaches in R&D, and encourage "efficient abandonment" by researchers. Learning in a research framework is studied, for instance, by Crawford and Shum (2005). Allen (1966) also examines individual R&D projects, and documents how information gathering differentially affects the progress of projects at different points in time.[4] However, we are aware of no other paper that estimates a dynamic learning of R&D decision-making at the level of the individual project. There is a parallel between our approach of using a dynamic model where expected benefits (citations) are observed to recover marginal costs, and the Pakes (1986) approach of estimating a model where marginal costs are observed in order to study the distribution of benefits.

We also contribute to the literature in organizational economics studying how the design of the institutional environment in IETF spurs successful research and development. The literature has analyzed the impact on innovation of subsidies to firms (Wallsten, 2000; Lerner, 1999), and how internal managerial practices and neighbors' R&D increase a firm productivity (Bloom et al., 2012, 2013). To our knowledge, we are the first to provide direct structural estimates of the impact of organization design on rate of learning and projects'

[4]Different approaches under consideration as solution to a problem see their breakthrough at different times; then differing durations of second phase. See Fig 2 on page 75.

efficient abandonment.

Our theoretical model is a mixture of two optimal stopping problems. The first phase is a variant of experimentation models using two-armed bandits.[5] We provide estimates of the success probability of the one-armed bandit (i.e., the rate of learning) in the context of internet standard development. A defining feature of our model is the second phase following this first phase of experimentation. In our framework, the prize of success is not deterministic, but is a function of the expected number of versions. Because of potential nonlinearities in the realized project values and opportunity costs, the expected continuation value upon breakthrough (i.e., success on the one-armed bandit) depends on the timing of breakthrough.

Our paper further relates to the literature on technology standardization in the IETF (Rysman and Simcoe, 2008; Simcoe, 2012). Our model is one of collaborative R&D and standardization. Alternative approaches have been taken by Ganglmair and Tarantino (2014), Hellmann and Perotti (2011), or Stein (2008). We also contribute to the empirical literature on this question of the importance of collaboration in economic activity (Wuchty et al., 2007). We are agnostic about the incentives within our author teams, but we find that larger research teams face lower opportunity costs. Moreover, collaboration within the IETF at large (through feedback sent in e-mails) is important because it increases the rate of learning.

The paper is structured as follows. In Section 2, we introduce the Internet

[5]Specifically, our model in this phase draws on the single-agent two-armed bandit problem in Heidhues et al. (2015), under the additional assumption that playing the safe arm is an absorbing state. Keller et al. (2005) study strategic experimentation in a continuous-time setting.

Engineering Task Force and provide details on the standardization process. In Section 3, we describe our data and provide simple descriptive and exploratory results. In Section 4, we present our two-phase Bayesian learning model of experimentation. In Section 5, we discuss the estimation procedure and identification strategy. In Section 6, we present the estimation results. In Section 7, we provide results on counterfactual simulations. In Section 8, we conclude.

2 The Internet Engineering Task Force

The IETF creates and maintains the technology standards used to run the Internet, such as the Transmission Control Protocol and Internet Protocol (TCP/IP) for routing packets. The organization was formed in 1986, and early members were primarily academic and government researchers. 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. It is now populated by researchers and engineers from public and private organizations (firms, universities, and other research centers).

The IETF has played a major role for the technological development of the Internet. Table 1 lists some of the more prominent standards certified by the organization. These include critical technologies tied to products in computer graphics, electronics, information technologies, and telecommunications. For instance, the Session Initiation Protocol (SIP) is the standard for the tech-

Table 1: Examples for IETF Internet Standards

	Description	RFC	Year
UTF-8	UTF-8, a transformation format of ISO10646	3629	2003
TIFF	Tag Image File Format (TIFF) – image/tiff MIME Sub-type Registration	3302	2002
SIP	Session Initiation Protocol	3261	2002
HTTP	Hypertext Transfer Protocol – HTTP/1.1	2616	1999
IPV6	Internet Protocol, Version 6 (IPv6) Specifiction	2460	1998
DHCP	Dynamic Host Configuration Protocol	2131	1997
MIME	Multipurpose Internet Mail Extensions MIME Part 1: Format of Internet Message Bodies	2045	1996
POP3	Post Office Protocol – Version 3	1939	1996
PPP	The Point-to-Point Protocol (PPP)	1661	1994
FTP	File Transfer Protocol	959	1985
TCP	Transmission Control Protocol	793	1981
IP	Internet Protocol	791	1981

nologies that enable internet service providers across the globe to offer VoIP ("Voice over IP") services. It supports video conferencing, instant messaging, file transfer, and online games, among others services.

A distinctive feature of the IETF is its transparency. It grants access to all intermediate and final versions of both published and abandoned projects on a public repository. This repository is managed and maintained by the organization, whose goal is to spur the participation of the members of the community. At the same time, the repository allows for the dissemination of the knowledge developed by the organization in the scientific community. The organization also provides access to an e-mail server on which much of the project-related communication between IETF members is published. Via e-mail discussion lists, members discuss the content of a proposal, provide feedback, and voice questions and concerns to be considered for a revised version.

2.1 The Standards Development Process

The following description of the IETF standards development process is based on Simcoe (2012). The process begins when participants identify a problem and form a working group (WG) to consider solutions. To prevent forum shopping and overlapping technical agendas, new working groups must be approved by an advisory board called the Internet Engineering Steering Group (IESG). Once a working group is formed, anyone can submit a proposal for a standard by posting it to the public repository. These proposals are referred to as "Internet Drafts" (ID). IDs are debated at triannual IETF meetings and on the e-mail discussion lists maintained by each working group. Much of the communication related to a project's revision process takes place via these e-mail discussion lists. IDs are continually revised, and, as a statutory rule, an unpublished ID expires after six months if the authors do not submit a new version.

For an ID to be published as a "Request for Comments" (RFC), the relevant working group 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. In practice, a working group chair decides whether consensus has been reached. If the working group chair declares a consensus, there is a "last call" for comments within the working group, 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 working group for resolution. If the IESG is satisfied that a consensus exists within the working group and sees no problem with the ID, it will be published as an RFC.

There are two types of RFCs. Standards-track RFCs define new protocols, which progress in maturity from Proposed Standard to Draft Standard and then finally to Internet Standard. Nonstandards-track RFCs are classified as Informational or Experimental. While standards and nonstandards go through an identical development and publication process, nonstandards do not receive an official endorsement and may not advance unless resubmitted as an ID for standards-track publication.[6]

3 Data and Descriptive Results

For our empirical analysis, we construct a novel project-level dataset of internet standard projects at the IETF. Our final sample, spanning the period of two decades, holds more than 16,000 completed projects. We have information on both successfully completed projects (*published* as RFC) and failed projects (*abandoned*), including the size of the project team, project-related communication by IETF members during a project's revision process, and the number of times each published project is cited in U.S. patents. The unique features of this dataset allow us to relate the development of a project, and its characteristics, to the information on whether it is published or abandoned. In this section, we describe the construction of the dataset and the main features

[6] We will exploit this feature of standard-track and nonstandard-track proposals in our estimation design.

of the final sample.

3.1 Sample Construction

For the universe of IETF projects, we download bibliographic information and all available version documents from the IETF repository. Individual versions are identified through an ID designation and a version number.[7] These designations may change over time, or different IDs are merged.[8] We use information provided by the IETF to link continuing IDs and thus construct projects as a series of IDs and versions.[9]

We restrict our sample to projects that were initiated in 1996 or later. We drop all projects that are active, that means, all projects that have not been completed and thus have not realized an outcome. Out of all active projects, 97% are initiated in 2010 or later. In order to avoid selecting projects based on outcome, we drop all completed projects that were initiated in those years. At last, we exclude a list of specialized projects.[10] This leaves us with a final sample of 16,268 completed projects, initiated in years 1996 through 2009 and

[7]For instance, the ID for the Hypertext Transfer Protocol (http) version 1.1 is `draft-ietf-http-v11-spec-rev`.

[8]For instance, the ID `draft-arkko-townsley-homenet-arch` is superseded by `draft-chown-homenet-arch`, which is later superseded by `draft-ietf-homenet-arch`, and eventually published as RFC 7368. We link these four IDs and treat them as a single project.

[9]When constructing these series of IDs and versions, the first available document in the first ID of a series is the first version of a project. For some (older) IDs initiated by individuals, early versions of a project are not available. We thus make the first available document our first version. We also encounter a total of 467 missing intermediate documents (accounting for about 0.5% of the total number of documents). We do account for missing documents and interpolate missing values when possible.

[10]The IETF repository also holds documents on projects associated with the Internet Research Task Force (IRTF), the Internet Architecture Board (IAB), Internet Engineering Steering Group (IESG), and the Internet Assigned Numbers Authority (IANA). We exclude these and focus on standards development within the IETF. We exclude projects designated as “best current practice”, “draft standard”, “historic”, or “internet standard”.

completed in years 1996 through 2015.

Completed projects are of one of two outcomes. We refer to projects that have expired or have been withdrawn by either the IETF or the submitter as *abandoned*. We refer to projects that are not abandoned, and thus successful, as *published* (as RFCs). In our sample, roughly 25% of all projects are published.

The ID of the first version of a project indicates whether the project was initiated within (or sanctioned by) a working group or an individual outside a working group.[11] Roughly 25% of all projects are initiated within a working group. We refer to the sample of these working group projects as the WG sample. Note that a considerable number of projects start off as individual projects but move into a working group, so that roughly 30% of all projects are completed working group projects.

For the size of the project team, we parse the text of the individual documents to obtain information on the authors of a given version.[12] We then construct a *team size* variable as the number of authors for a given version. We use the team size on the initial version as our project-specific value.[13]

To capture the extent of involvement of the IETF community at large (reflecting community attention), we construct a variable of project-related communication. We exploit the following feature of the IETF process. Each

[11] IDs initiated within a working group start with `draft-ietf-`; "individual" IDs start with `draft-[...]-`.

[12] We use Jari Arkko's Perl script which can be downloaded at http://www.arkko.com/tools/docstats. We manually collect information on authors in about 800 documents for which the Perl script does not return any information.

[13] The team size is fairly consistent over time. The mean number of authors on the initial version is 2.32 (std. dev.: 1.89) whereas the number of authors on the final version is slightly higher at 2.47 (std. dev.: 1.86).

version of a project is announced through an e-mail to members of the IETF, and a large part of the ensuing discussion of a version is via e-mail discussion lists. Using the ID designation and version number, we match a given version of a project with all e-mail messages sent in response to that version (i.e., all e-mail messages sent between a version t and the next version $t+1$). The sum of all e-mail messages divided by the number of versions is the measure of project-related communication.

We construct a measure of patent citations to capture the *value* of a project. Using the full text of U.S. patents from 1976 through 2015,[14] we count the number of patents that cite a given IETF project.[15] We find that a considerable number of U.S. patents cite IETF projects before these are published as RFCs. For our value measure, we use patent citations only for published RFCs. We use the predicted log of citations (with a base year of 2005) as the realized value of an RFC, denoted by $\hat{\pi}(t)$.

3.2 Descriptive Results

Table 2 provides summary statistics for our main variables for both the full sample and the WG sample. We also break down the numbers by whether a project is *on* or *off* the standard track.

We see considerable variation of our key variables between as well as within samples. In the full sample, 24.5% of all projects (15.5% on the standard track)

[14] The PatentsView project provides flat-format files of full-text patents at http://www.patentsview.org/download/.

[15] We take three different approaches to search for IETF project citations in patents: (1) search by RFC number; (2) search by ID; (3) search by project title (only long titles, in combination with queries "internet draft", "internet standard", "IETF", or "Internet Engineering Task Force."

Table 2: Sample Statistics

	Full Sample			WG Sample		
	All	On Track	Off Track	All	On Track	Off Track
Projects	16268	14549	1719	3982	3201	781
Versions/Project	3.67	3.3	6.81	5.67	5.42	6.7
(Std.Dev.)	(4.17)	(3.94)	(4.68)	(4.87)	(4.96)	(4.35)
% Projects in WG (first)	24.5%	22%	45.4%			
% Projects in WG (last)	30.2%	26.6%	60.2%			
% Projects Published	24.5%	15.5%	100%	55.5%	44.7%	100%
Length (in Words)	575.7	557.4	650.8	652.2	638.9	696.5
(Std.Dev.)	(349.9)	(283.5)	(536.8)	(459.6)	(334.5)	(734.4)
Projects per Year	1162	1039.21	122.79	284.43	228.64	55.79
1996–2000 (p.a.)	917.8	807.8	110	366.4	307.6	58.8
2001–2005 (p.a.)	1607.2	1440.4	166.8	424.4	334.8	89.6
2006–2009 (p.a.)	1446.75	1263.25	183.5	250.25	190.5	59.75
E-mail/Version	4.87	4.8	5.42	5.27	5.07	6.11
(Std.Dev.)	(7.53)	(7.66)	(6.27)	(6.55)	(6.58)	(6.33)
Team Size	2.32	2.3	2.48	2.5	2.44	2.76
(Std.Dev.)	(1.89)	(1.89)	(1.89)	(2.42)	(2.49)	(2.1)
1 Author	40.9%	41.6%	35.8%	39.1%	41%	31.6%
2 Authors	26.6%	26.4%	28.8%	26.6%	26.6%	26.9%
3-4 Authors	23.7%	23.6%	24.3%	23.3%	22.5%	26.8%
5+ Authors	8.8%	8.5%	11.1%	10.9%	10%	14.7%
Citations (RFC)	9.85	12.11	6.92	13.94	16.4	9.48
(Std.Dev.)	(32.22)	(39.35)	(19.02)	(41.04)	(47.74)	(23.92)
1996–2000 (p.a.)	1.24	1.52	0.89	1.48	1.66	1.14
2001–2005 (p.a.)	0.83	1.1	0.46	1.04	1.34	0.5
2006–2009 (p.a.)	0.39	0.48	0.27	0.34	0.43	0.17
1 Author	7.89	10.26	4.53	12.4	15.06	6.91
2 Authors	9.16	10.7	7.38	12.63	13.88	10.41
3–4 Authors	11.05	12.96	8.48	14.53	16.94	10.3
5+ Authors	15.73	21.44	9.23	19.59	25.51	10.89

Standard deviations in parentheses.

are published, whereas the publication rate in the WG sample is 55.5% (44.7% on the standard track). Working group projects have on average more versions (5.67 vs. 3.67), and we observe longer processes off-track than on-track (6.81 vs. 3.3). The conditional probability for each of the two outcomes varies with the number of version. We see this in the lower-left panel in Figure 1. It depicts the probabilities of publication (solid line) and abandonment (dashed line) as function of a project version. The probability of publication exhibits an increasing pattern and eventually levels off, reaching a maximum of about 19% at version 17. The probability of abandonment, on the other hand, decreases with a project version number. Specifically, 40% of the projects are abandoned after the initial version. These results suggest that, while members of the community learn fairly fast whether a project should be abandoned, it takes a considerably larger number of versions for the project to be ready for publication.

We further document that this revision process has real effects. A project's duration is associated with the length of a project's specification (in terms of unique words) as a measure for content: projects have more content off-track than on-track (650.8 vs. 557.4) and when initiated within working groups (652.2 vs. 575.7). We illustrate the effect of the number of versions on content in the upper-left panel of Figure 1. It plots the text distance of a given version T from the initial version.[16] This text distance of the average project in our sample increases at a decreasing rate, a pattern that suggests that the average revision process of a project develops "away" from the initial version, but incremental changes decrease over version-time. A priori, this is

[16]To construct this measure, we use techniques borrowed from text-based analysis. A text document (i.e., a version document) is represented as a vector of word frequencies. A common measure of textual distance is the cosine distance:

$$1 - \frac{x_T \cdot x_1}{\| x_T \| \| x_1 \|}$$

where x_T is the vector of word frequencies for version T and x_1 the vector for the initial version.

Figure 1: IETF Project Development

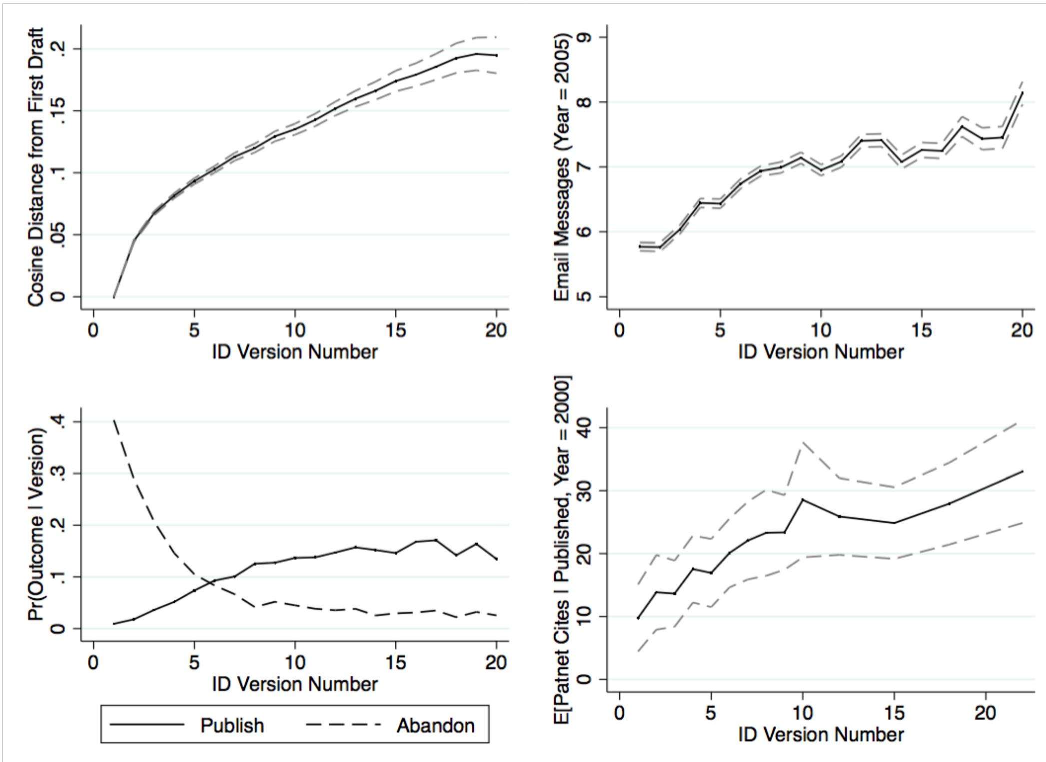

Top-left: text distance of version T from initial version 1 for varying ID Version Number. Top-right: predicted average number of e-mail messages per version for base year 2005. Bottom-left: probabilities for publication (solid line) and abandonment (dashed line) as function of project (ID Version) number. Bottom-right: predicted patent citations (of RFCs) for base year 2005.

not obvious because the pattern could exhibit strong non-monotonicities and thus reflect the presence of disagreement among members of the committee. The documented monotonicities motivate our later assumption of cooperative decisions within author teams.

Off-track projects and WG sample projects also attract more project-related communication (E-mail/Version), which intensifies as the project undergoes more revisions. The top-right panel of Figure 1 depicts the average number of e-mail messages (per version) exchanged over the course of a project.

We plot the in-sample prediction of the expected number of e-mail messages exchanged in response to an average version for the 2005 base year.[17] It shows that the number of e-mail messages per version increases with the version number, suggesting that the community becomes more active as the number of versions increases.

The figures in Table 2 further suggest that published projects in the WG sample receive (on average) more citations (13.94 vs. 9.85); as do on-track projects (12.11 vs. 6.92). Also, RFCs published earlier (1996–2000) on average receive more citations per annum than RFCs cited in later years (1.24 vs. 0.83 in 2001–2005 and 0.39 in 2006–2009). We see similar variation when breaking down citations by the number of authors on a project. Projects with more authors receive more citations (15.73 for RFCs with 5+ authors vs. 7.89 for RFCs with 1 author). The bottom-right panel of Figure 1 plots the in-sample prediction of the expected number of citations received by an RFC for the 2005 base year. On average, RFCs with more versions receive more citations. This finding implies a positive and strong correlation between the number of versions of a project and its value as captured by patent citations.

Table 2 and Figure 1 have two main take-aways. First, there is ample heterogeneity to exploit in the empirical analysis within and between samples. We use this heterogeneity in an extension of our model where we allow the model parameters to vary across subsamples. Second, the statistical features of the projects in the WG sample suggest that there is potential self-selection of better (i.e., more likely to be publishable) projects into this subsample. We

[17] The choice of alternative years leads to analogous results.

Table 3: Publication, Learning, and Productivity

Specification	OLS	Poisson	Poisson
Outcome	Published	Versions	Citations
log(Versions)	0.25		0.75
	[0.00]**		[0.08]**
log(E-mail/Version)	0.02	-0.08	0.29
	[0.00]**	[0.01]**	[0.04]**
Published * log(E-mail/Version)		0.12	0.09
		[0.01]**	[0.07]
Published		0.95	1.64
		[0.02]**	[0.09]**
No E-mail	0.03	-0.58	-0.10
	[0.01]**	[0.02]**	[0.16]
WG Project	0.26	0.48	0.15
	[0.01]**	[0.02]**	[0.07]*
Cohort Effects	Y	Y	N
Publication Year Effects	N	N	Y
Observations	16,271	16,271	16,271

* 5% significance; ** 1% significance.

document this by estimating our model both on the full sample and the WG sample.

3.3 Reduced Form Regression Results

Table 3 presents some exploratory regression results that capture some of the relationships depicted in Figure 1, and help to motivate the model we develop below. All of these regressions are cross sectional, based on a sample consisting of the last version of every project in the IETF sample.

The first column in Table 3 presents estimates from a linear probability model of publication. It shows that there is a strong association between the number of version of a project and the probability of publication. Doubling the number of versions increases the probability of RFC publication by around 17 percentage points. Working group projects are also 26 percentage points more likely to be published. There is a positive and statistically significant, but economically much smaller relationship between the volume of e-mail linked to a project and its likelihood of success. A one standard deviation increase in e-mail messages per version increases the publication probability by roughly 2.5 percentage points.

The second column in Table 3 examines the link between communication and revisions, and motivates the type of Bayesian learning we model below. As in the first column, we see a very large and strong association between publication and the number of version. However, in this regression we also observe that the e-mail per revision variable is negatively correlated with the number of version for unpublished projects, and positively associated with versions for published projects (after controlling for the "low end" projects that receive no e-mail and fail very quickly). This difference-in-difference results suggest a link between communication and learning. In particular, more active communication seems to lead to "fast failure" for unpublished projects and more versions for those eventually published. Our model below will incorporate the idea that faster learning leads projects that have not experienced a "breakthrough" to drop out more quickly.

The IETF becomes more efficient if projects that are unlikely to be published drop out more quickly. However, failing fast can produce a trade-off if developers give up too soon on ideas that might become successful given more versions. One way to look at the overall productivity of the IETF is to examine a citation production function that treats both version and communication as inputs. That is what we do in the third column of Table 3. Although we do observe some citations to unpublished projects, the model shows a very large (roughly 500 percent) increase in cites to published RFCs. We also see a very strong positive association between both versions and e-mail communication and the expected patent citations to a project.

One way to deal with the endogenous "input" of versions is to model the process of deciding to continue working on a project. If the authors stopped as soon as they knew a project could be published, we might expect failures to take longer than successes. But instead we observed the opposite – projects published as RFCs go through more revisions, and the number of revisions is correlated with the number of follow-on cites. In the next section, we develop a model of Bayesian learning in R&D that captures each of these features of the data.

4 A Bayesian Learning Model of Experimentation in Internet Standards Development

4.1 Overview

For our model of internet standards development, we consider a process in which a team of researchers jointly develops a project. The team is endowed with an initial version of a project of unknown quality and can revise its

Figure 2: Stylized R&D

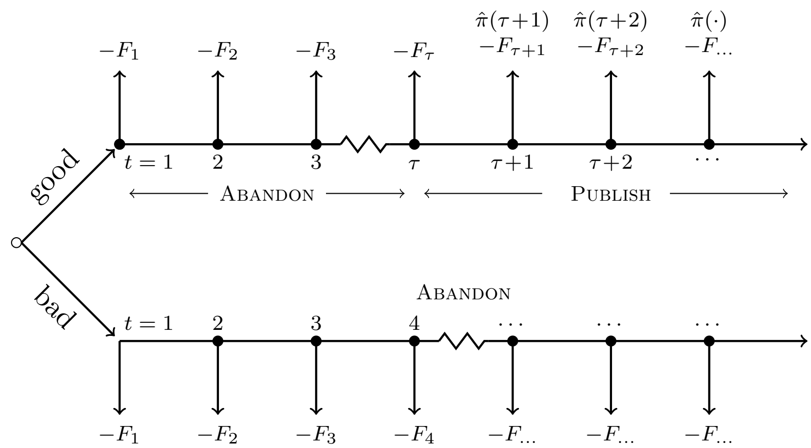

specifications both to learn its type and to increase its (potential) value.[18] We assume that the project quality can be either good or bad, and the project generates value (i.e., materializes its potential value) for the team only if it is of the good type, and the team has observed the type before the process ends. Each costly version of the project increases the potential value and allows the team to run a new experiment. In other words, not yet having learned the type, the team can "experiment" by submitting a new version to realize whether the project type is good. We refer to the realization of the good type as a *breakthrough* that changes the status of the process. We assume that, conditional on a good type, this experimentation process is successful (and a breakthrough occurs) with constant probability.

Figure 2 provides a stylized depiction of this process. The realized potential value of a version t is denoted by $\hat{\pi}(t)$. Suppose the costs of a version n is $F(n)$. Then cumulative (non-stochastic) costs of version t are denoted by

[18]We do not consider the team's entry decision by assuming that the initial version comes at no cost.

$F_t = \sum_{k=0}^{t-1} F(k)$, with $F(0) = F_1 = 0$.[19] Bad projects never experience a breakthrough and will never realize their potential value. Good projects may or may not experience a breakthrough in a period $t = \tau$. The team decides in each $t \geq 1$ whether to submit a revised version of the project specifications or stop the process. After a breakthrough, the decision to continue or stop is a simple comparison of the incremental value of a version and the costs of the version. Before a breakthrough, the team must form beliefs about the type of the project. The more failed experiments the team has observed (without a breakthrough), the more pessimistic it will be that the type is good. On the other hand, the more versions the team has submitted, the higher is the value to be realized when a breakthrough occurs. Thus, if it stops before a breakthrough, $t < \tau$, the team loses the opportunity to harvest the potential value of the project while having incurred the cumulative costs of all version $t' = 1, \ldots, t-1$. If it stops after a breakthrough, its payoffs are the potential value net of the cumulative costs of all prior versions.

Before we introduce more notation to formalize these ideas, we find it useful to relate this general innovation process to the procedures within the IETF. In this context, the project is an internet standard to-be-developed by a team of engineers. The team has an initial status of the standard. The goal is to develop specifications that are endorsed by the IETF and published as an RFC. We refer to the anticipated endorsement as a breakthrough. For example, the breakthrough happens if (a qualified majority of) participants in a working group agree on the commercial interest, or the technical merits, of the tech-

[19] For a version t, the team must incur revised versions and incur costs $F(t')$ in all $t' = 1, \ldots, t-1$.

nology under discussion. A project that has experienced a breakthrough will eventually be published as an RFC once the team decides to stop the process.

In our empirical setting we assume that both the team and the econometrician observe the value of $\hat{\pi}(T)$, for all T. The probability that a project is of the good type, denoted by p, and the probability that a breakthrough occurs in $t \leq \tau$, denoted by b, are known by team of researchers, but unknown to the econometrician. At the same time, while the team observes when a breakthrough has occured, the econometrician only observes whether but not when it takes place during a project. The identification of the parameters for the project type and the breakthrough probabilities is our main task in the estimation of this model of experimentation. Moreover, a major challenge to estimation is the unobservability of the exact timing of the breakthrough.

4.2 Model Setup

In what follows, we provide a more technical account of our model. In $t = 0$, a team of risk neutral agents initiates a project of type $\theta \in \{\text{good}, \text{bad}\}$. We assume this initial version of the project comes at zero costs; we therefore ignore the team's entry decision. The team initially does not know the type of the process, but has prior beliefs $p = \Pr(\theta = \text{good})$ that the project is good, with $0 < p < 1$ and $\Pr(\theta = \text{bad}) = 1 - p$.

The project type is payoff-relevant insofar as only good projects whose type has been realized generate value. The good type is realized when the team learns that the project is good, that means, when a *breakthrough* occurs. Let b denote the per-period probability of a breakthrough. Given a breakthrough has not occurred in any $t' < t$, a breakthrough can occur in t only if the following two conditions hold:

1. *The project is of the good type.* This implies $\Pr(\text{breakthrough in } t|\theta = \text{good}) = b$ and $\Pr(\text{breakthrough in } t|\theta = \text{bad}) = 0$.

2. *The team has submitted a version of the project.* This implies that learning requires experimentation in the form of a revision (a new status).

Once a breakthrough occurs in τ, the good type is realized. In the context of the IETF this means that the content of the version is endorsed and eventually published. Any additional version at this point will improve the content of the project but will not affect the status of the project. We use σ_t to denote this status in t:

$$\sigma_t = \begin{cases} 0 & \text{for } t \leq \tau \qquad [\text{“pre-breakthrough phase”}] \\ 1 & \text{for } t > \tau \qquad [\text{“post-breakthrough phase”}] \end{cases} \tag{1}$$

In each t, the team forms beliefs $\hat{p}(t|\sigma_t)$ about the type of the project. When a breakthrough has occurred and the status is $\sigma_t = 1$, then $\hat{p}(t|1) = 1$ for all $t > \tau$. If, instead, in a given t a breakthrough has not yet occurred, then posterior beliefs are formed by Bayes' rule:

$$\hat{p}(t|0) = \frac{p\,(1-b)^t}{(1-p) + p\,(1-b)^t} \tag{2}$$

with the prior $p = \hat{p}(0|0)$. With probability $1 - b$ this first experiment fails and the team updates beliefs to $p(1|0) < p$ for its decision in $t = 1$.[20] As long

[20]We provide below a more detailed description of the sequence of events within each

as experimentation fails and breakthrough does not occur, we have $\hat{p}(t|0) < p$ for all $t \leq \tau$.

We use $\pi(t|\sigma_t)$ to denote the ex-post realized value of a successful project as function of the number of versions, t, and the status of the process, σ_t. Recall that unless the good type of the project has been realized (and the status of the process is $\sigma_t = 1$), the project generates no value. The ex-post value of a project can thus be summarized as

$$\pi(t|\sigma_t) = \begin{cases} \hat{\pi}(t) & \text{if } \sigma_t = 1 \\ 0 & \text{if otherwise} \end{cases} \tag{3}$$

and we assume that $\hat{\pi}(t)$ is non-decreasing in t.

In each $t \geq 1$, the team cooperatively decides to continue or stop. We assume the team does not discount future payoffs. The sequence of steps in each period t is as follows:

$t.1$: Given the outcome of previous rounds' experimentations, the team updates its beliefs about the type of the project. If $t > \tau$ and $\sigma_t = 1$, posterior beliefs that the project is good are $\hat{p}(t|1) = 1$. If a breakthrough has not been observed in a previous round and $\sigma_t = 0$, the posterior beliefs that the project is good are $\hat{p}(t|0)$ according to the expression in equation (2).

$t.2$: The team observes a cost shock ε_t. The incremental cost of a new version t are $F(t) + \varepsilon_t$ with $F(0) = 0 = \varepsilon_0$ for the initial version. We assume

period t.

that the non-stochastic cost component is strictly positive, $F(t) > 0$ for all $t > 0$. Upon observing the cost shock ε_t, the team decides to continue by submitting a new version of the project.

$t.3$: If in $t.2$, for $\sigma_t = 0$, the team decides to continue, it observes the outcome from experimentation. If experimentation is successful (a breakthrough occurs), then $\tau = t$ so that $\sigma_{t+1} = 1$, and the team moves to the post-breakthrough phase; otherwise, $\sigma_{t+1} = 0$ and the team stays on the pre-breakthrough phase.

If the team stops, the project is abandoned. In the model, abandonment is an absorbing state. This assumption reflects the institutional features of the IETF, which imposes a six-month rule after which, if no new version is submitted, projects expire. Alternatively, this assumption reflects the presence of depreciation of knowledge in the development of a project.

When the team decides in $t.2$ whether to continue or stop, it compares the payoffs from stopping in t with the expected value of another version, net of the costs. Let $EV(t|\sigma_t) = E(V(t+1)|\sigma_t) - \pi(t|\sigma_t)$ be the expected (option) value of continue relative to stop, where $\pi(t|\sigma_t)$ denotes the value when the team stops and $E(V(t+1)|\sigma_t)$ denotes the expected value of another version given the current status σ_t of the project. We characterize these value functions in greater detail below. The team continues in t if $EV(t|\sigma_t) \geq F(t) + \varepsilon_t$ and stops otherwise. We assume that $F(t) > 0$ for all $t > 0$. The continuation decision can be rewritten as:

$$\begin{aligned}\varepsilon_t \leq \bar{\varepsilon}_t^{\sigma_t} &:= EV(t|\sigma_t) - F(t) \\ &= E(V(t+1)|\sigma_t) - \pi(t|\sigma_t) - F(t). \end{aligned} \tag{4}$$

The team continues in t as long as the cost shock does not exceed the critical threshold $\bar{\varepsilon}_t^{\sigma_t}$. From an ex-ante point of view, this means that the team continues in period t with status σ_t with probability

$$G^{\sigma_t}(t) = \Pr(\varepsilon_t \leq \bar{\varepsilon}_t^{\sigma_t}). \tag{5}$$

4.3 Expected Payoffs

The goal of the team is to formulate a contingent plan of actions that maximizes its expected payoffs. To characterize these payoffs, it helps to first characterize the probabilities of the two possible outcomes, given that the team stops in a period T. This T is the final number of versions.[21] If, in this final period T, the good type has been realized so that $\sigma_T = 1$, then the project is a success and *published* with a value of $\pi(T|1) = \hat{\pi}(T)$. If, in T, the good type has not been realized (either because the project is of the bad type or a breakthrough has not occurred for a good project), then the project is a failure and abandoned with value $\pi(T|0) = 0$.

In Figure 3, we illustrate the decision trees for a good and a bad project,

[21] At the outside of the game, there is one initial version. Once the team stops in $t = 1$, it stops with one version but has not occurred any costs. If it stops in $t = 2$, there are two versions (after continuing in $t = 1$), it incurs costs of continuing in $t = 1$.

as determined in $t = 0$. The branch for the good type (in panel 3a) is reached with probability p, the branch for the bad type (in panel 3b) is reached with probability $1 - p$.

Consider the timeline of good projects in Figure 3a. We assume an initial version (first experimentation) exists in $t = 0$. The phase is then such that $\sigma_0 = 0$. With probability b, this experimentation is successful, and the project is in the post-breakthrough phase (with status $\sigma_1 = 1$) in $t = 1$. With probability $1-b$, experimentation in $t = 0$ is not successful and the status remains $\sigma_1 = 0$. Moving forward, in $t = 1$ with status σ_1, the team now decides to continue or stop. From the viewpoint of $t = 0$, the team continues with probability $G^{\sigma_1}(1)$ and stops with probability $1 - G^{\sigma_1}(1)$. If the team stop, the project is published (when in the post-breakthrough phase) or abandoned (otherwise). If the team continues, the post-breakthrough phase moves to $t = 2$. The pre-breakthrough phase moves to $t = 2$ (with $\sigma_2 = 0$) with probability $1 - b$; it moves to the post-breakthrough phase (with $\sigma_2 = 1$) with probability b. The game proceeds until the team decides to stop. We denote this last period in which decisions are made by T.

Consider the timeline for bad projects in Figure 3b. Because a breakthrough cannot occur for bad projects, the status of the project is $\sigma_t = 0$ for all t.[22] In $t = 1$, the team continues with probability $G^0(1)$ and stops with probability $1 - G^0(1)$. If the team stops, the project is abandoned. If it continues, the pre-breakthrough phase moves to $t = 2$. The game proceeds until the the team decides to stop.

[22]To be precise, there is a post-breakthrough phase with $\sigma_t = 1$ in this timeline, but it is reached with probability zero in each t.

Figure 3: Timeline for Project Types

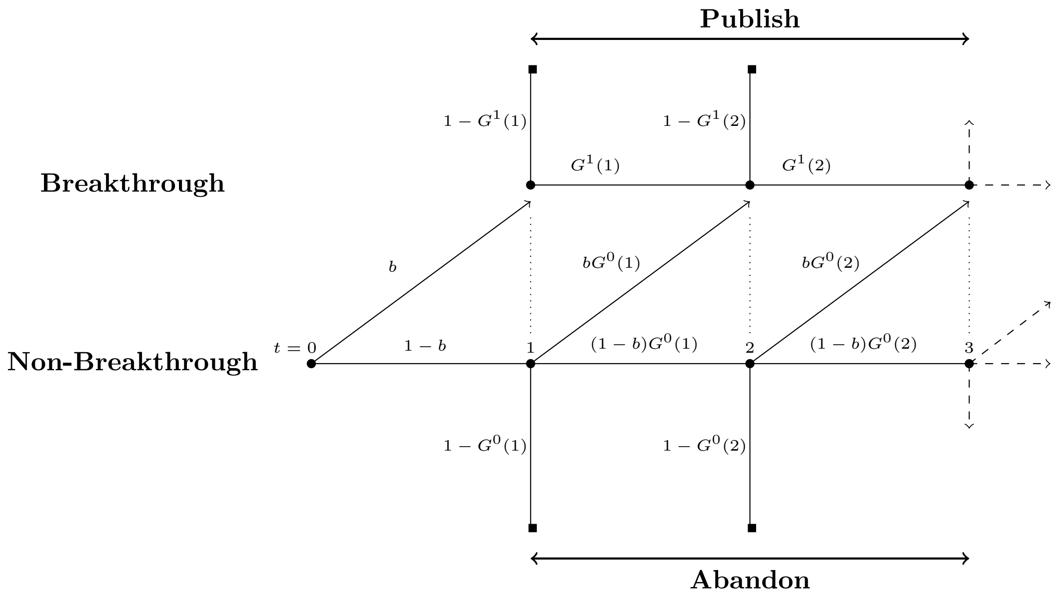

(a) Timeline for Good Projects

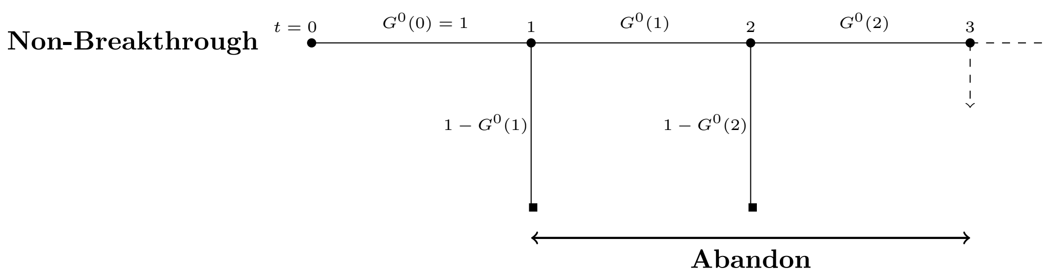

(b) Timeline for Bad Projects

We denote the probability that a project is abandoned in T by $\Phi_A(T)$. The project is abandoned in T if it is bad or a breakthrough has not occurred for a good project, and the team has continued until T, with sufficiently low cost shocks ε_t in all $t \leq T-1$ and a sufficiently high cost shock in $t = T$. The respective critical thresholds for the costs shocks are $\bar{\varepsilon}_t^0$ as defined in equation (4) and continuation probabilities are $G^0(t)$ as defined in equation (5). The probability of abandonment in T is:

$$\Phi_A(T) \equiv \left[(1-p) + p\,(1-b)^T\right]\left(1 - G^0(T)\right)\prod_{k=0}^{T-1} G^0(k). \tag{6}$$

The expected sum of the incurred cost shocks for an abandoned project is:

$$E_A(T) \;\equiv\; \sum_{k=0}^{T-1} E(\varepsilon_k | \varepsilon_k \leq \bar{\varepsilon}_k^0). \tag{7}$$

We further denote the probability that a project is published in T by $\Phi_P(T)$. A project is published only if it has a breakthrough, and a breakthrough can occur only for good projects (with probability p). The team stops the process in T when the cost shock ε_T is too high (with ex-ante probability $1 - G^1(T)$). The probability that the project reaches this final T depends on the continuation decisions in $t = 1, \ldots, T-1$, which are conditional on the phase. Suppose the breakthrough occurs in some period $0 \leq \tau \leq T-1$. This implies period τ is reached without a breakthrough with probability $(1-b)^\tau$ and a breakthrough occurs with probability b.[23] Moreover, period τ is reached if all cost shocks are sufficiently low, with probability $\prod_{j=0}^{\tau} G^0(j)$. With $\tau = t$, the decisions in $t \geq \tau+1$ are in the post-breakthrough phase. Once a breakthrough has occurred, the team continues in a given $t > \tau$ with probability $G^1(t)$. The final period T is thus reached with probability $b\,(1-b)^\tau \prod_{j=0}^{\tau} G^0(j) \prod_{k=\tau+1}^{T-1} G^1(k)$ for a given τ. Summing up over all possible

[23] To see this, $t = 1 = \tau$ is reached without a breakthrough if the initial version does realize the good type, with probability $1 - b$. See Figure 3a.

$\tau < T$, we obtain the expression for $\Phi_P(T)$:

$$\Phi_P(T) \equiv p\left(1 - G^1(T)\right) \sum_{\tau=0}^{T-1} \left(b\,(1-b)^\tau \prod_{j=0}^{\tau} G^0(j) \prod_{k=\tau+1}^{T-1} G^1(k) \right). \tag{8}$$

The expected sum of incurred cost shocks for a published project is:

$$E_P(T) \;\; \equiv \;\; \sum_{\tau=0}^{T-1} b\,(1-b)^\tau \left(\sum_{j=0}^{\tau} E(\varepsilon_j | \varepsilon_j \leq \bar{\varepsilon}_j^0) + \sum_{k=\tau+1}^{T-1} E(\varepsilon_k | \varepsilon_k \leq \bar{\varepsilon}_k^1) \right). \tag{9}$$

For the expected net value of a project, we obtain:

$$\sum_{T=1}^{\infty} \left(\Phi_P(T)\,[\hat{\pi}(T) - F_T - E_P(T)] - \Phi_A(T)\,[F_T + E_A(T)] \right) \tag{10}$$

with $\sum_{T=1}^{\infty} (\Phi_A(T) + \Phi_P(T)) = 1$ and $F_T = \sum_{k=0}^{T-1} F(k)$.

4.4 Recursive Characterization

We solve for the team's decision problem recursively, under the assumption that the team's horizon is finite. Specifically, we assume that $\overline{T}$ is the maximum number of version after which the market value of the potential standard is zero. We then characterize the team's value function $V(\cdot)$ in more detail.

In a period t, if the team decides to stop, its payoffs are $\pi(t|\sigma_t)$. Alternatively, the team can decide to run a new experiment, pay the cost and submit a new version. The value of this option is equal to the expected value of the project in $t+1$, conditional on the current information. Formally, the team value function is

$$V(t) = \max\{\pi(t|\sigma_t), E(V(t+1)|\sigma_t) - F(t) - \varepsilon_t\}, \tag{11}$$

for all $t = 1, ..., \overline{T}$ (with no decision $t = 0$). The team's choice variable is the decision to continue. The observable state variable is given by the number of versions, t; the unobservable state variable is given by the posterior beliefs $\hat{p}(t|0)$ in t, which evolve according to a first-order Markov process.[24]

This property has two implications. First, the team's decision problem is non-stationary: The problem in (11) depends on whether $t > \tau$ (post-breakthrough phase) or $t \leq \tau$ (pre-breakthrough phase). Second, as we shall see when deriving the likelihood function implied by our model, non-stationarity introduces serial correlation in the controlled stochastic process generating the value functions $\{V(t)\}_{t=1}^{\overline{T}}$.[25]

Before proceeding with the characterization of the team's dynamic optimization problem in the two phases, recall that the team solves the stopping problem under the assumption that *stop* is an absorbing state (so that, if a new version of the project is not submitted, it is understood that the project is terminated). Thus, the goal is to determine the version $T \leq \overline{T}$ in which the team decides to stop the revision process. In what follows, we begin with the post-breakthrough phase.

[24]Conditional on $\sigma_t = 0$, the probability that the project is good in t depends on the status in $t - 1$.

[25]We refer to the stochastic process generating $\{V(t)\}_{t=1}^{\overline{T}}$ as "controlled" because, although it is inherently random, it is also affected by the team's decision to continue.

Post-Breakthrough Phase Assume a breakthrough has taken place in a $\tau < t$, so that $\sigma_t = 1$. In a stage t of the post-breakthrough phase, the value function is

$$V(t) = \max\{\hat{\pi}(t), V(t+1) - F(t) - \varepsilon_t\}, \tag{12}$$

for all $t = 1, ..., \overline{T}$. Given $\sigma_t = 1$, we have $\pi(t|0) = \hat{\pi}(t)$ and $E(V(t+1)|1) = V(t+1)$. The solution for the sequence of $\{V(t)\}_{t=1}^{\overline{T}}$ in the post-breakthrough phase can be obtained starting with the terminal period, $\overline{T}$. The team stops and obtains payoffs of $\pi(\overline{T}|1) = \hat{\pi}(\overline{T}) > 0$. In $\overline{T} - 1$, value function becomes $V(\overline{T}-1) = \max\{\hat{\pi}(\overline{T}-1), \hat{\pi}(\overline{T}) - F(\overline{T}-1) - \varepsilon_{\overline{T}-1}\}$ and the team continues if $\hat{\pi}(\overline{T}) - F(\overline{T}-1) - \varepsilon_{\overline{T}-1} \geq \hat{\pi}(\overline{T}-1)$ and stops otherwise. The team then solves this problem backwards through $t = 1$.

Pre-Breakthrough Phase Assume a breakthrough has not taken place in any $t' < t$, so that $\sigma_t = 0$. In a stage t of the pre-breakthrough phase, the value function is

$$V(t) = \max\{0, E(V(t+1)|0) - F(t) - \varepsilon_t\}, \tag{13}$$

for all $t = 1, ..., \overline{T}$. Given $\sigma_t = 0$, we have $\pi(t|0) = 0$ and

$$E(V(t+1)|0) = b\hat{p}(t|0)E(V(t+2)|1) + (1 - b\hat{p}(t|0))\, E(V(t+2)|0) \tag{14}$$

When the team continues in t, it incurs costs $F(t)+\varepsilon_t$ and expects continuation payoffs $E(V(t+1)|0)$ as in (14), where a breakthrough occurs with probability $b\hat{p}(t|0)$ and does not occur with probability $1 - b\hat{p}(t|0)$.

We solve for the sequence of $\{V(t)\}_{t=1}^{\overline{T}}$ in the pre-breakthrough phase by starting with the terminal period, $\overline{T}$. The process ends and the team's are $\pi(\overline{T}|0) = 0$. In $\overline{T} - 1$, the value function becomes

$$V(\overline{T}-1) = \max\{0, b\hat{p}(\overline{T}-1|0)\hat{\pi}(\overline{T}) - F(\overline{T}-1) - \varepsilon_{\overline{T}-1}\},$$

and the team continues if the expected payoffs are nonnegative. The team then solves this problem backwards through $t = 1$.

4.5 Likelihood Function

For the likelihood function, we begin with the likelihood of publication. Because the project stops in T, the last period in which a breakthrough can occur is $T-1$. For projects (with T versions) published as RFCs, we know that there was a breakthrough but do not know for which version t. This implies that the likelihood of publication in T depends on the status of the project in T, σ_T, but also on the status of the project in all $t < T$ (and, in particular, on the value of τ). This property of the model introduces a simple form of serial correlation across the team's continuation decisions into our likelihood function.

In order to account for the fact that the breakthrough might occur at any $\tau = 0, ..., T-1$ (and is observed by the team at the beginning of any

$\tau+1$), it is helpful to define a function $\rho(\tau, T)$ as the probability of observing a breakthrough in period $t=\tau$ and publishing an RFC in $t=T$. This probability is equal to

$$\rho(\tau, T)=b\,(1-b)^{\tau}\left(1-G^{1}(T)\right) \prod_{j=0}^{\tau} G^{0}(j) \prod_{k=\tau+1}^{T-1} G^{1}(k). \tag{15}$$

Summing over all possible periods in which a breakthrough can occur, we can write the likelihood of publication in T as

$$p \sum_{\tau=0}^{T-1} \rho(\tau, T).$$

The log-likelihood for publication in a given T is equal to

$$LL_{\text{publish}}(T|\boldsymbol{\sigma_T}, b, p, \boldsymbol{F})=\log(p)+\log\left(\sum_{\tau=0}^{T-1} \rho(\tau, T)\right), \tag{16}$$

with $\boldsymbol{\sigma_T}=(\sigma_0, \sigma_1, \ldots, \sigma_T)$, and $\boldsymbol{F}$ denotes a vector of cost parameters.

The likelihood of abandonment in T is

$$(1-p) \prod_{k=0}^{T-1} G^{0}(k)\left(1-G^{0}(T)\right)+p\,(1-b)^{T} \prod_{k=0}^{T-1} G^{0}(k)\left(1-G^{0}(T)\right).$$

The first term refers to bad projects for which a breakthrough is not possible (status $\sigma_t=0$ for all t), whereas the second term refers to good projects that do not get a breakthrough (status $\sigma_t=0$ for all $t \leq T$). The log-likelihood for abandonment in a given T is equal to

$$\begin{aligned} LL_{\text{abandon}}(T|\boldsymbol{\sigma_T}, b, p, \boldsymbol{F}) &= \sum_{k=0}^{T-1} \log(G^0(k)) + \log(1 - G^0(T)) \\ &\quad + \log\left((1-p) + p\,(1-b)^T\right). \end{aligned} \tag{17}$$

We can now write the log-likelihood of the data. Let i denote a project and $\mathcal{I}$ the set of all projects in our sample. Project i ends in $t = \widetilde{T}_i$ with outcome $a_i \in \{\text{abandon}, \text{publish}\}$. The log-likelihood of the data (given our parameters and the vector of past statuses) is equal to:

$$LL(\boldsymbol{\sigma}, b, p, \boldsymbol{F}) = \sum_{i \in \mathcal{I}} LL_{a_i}(\widetilde{T}_i | \boldsymbol{\sigma_{\widetilde{T}_i}}, b, p, \boldsymbol{F}). \tag{18}$$

5 Estimation

5.1 Empirical Approach

We estimate the described dynamic decision problem in discrete time, were time is in version-time, t. This means, the duration of a project is equal to the number of versions, T. We maximize the likelihood function $LL(\boldsymbol{\sigma}, b, p, \boldsymbol{F})$ over $(b, p, \boldsymbol{F})$.

For the non-stochastic cost component, we assume that $F(t)$ is quadratic in version-time with $F(t) = C_0 + C_1 t + C_2 t^2$ and $\boldsymbol{F} = (C_0, C_1, C_2)$. For the stochastic cost component we assume that the cost shocks ε_t are independent draws from a Type I logistic distribution, $\varepsilon_t \sim \text{Logistic}(0, 1)$, $t \geq 1$. We denote the CDF of this distribution by G, with

$$G(\varepsilon) = \frac{1}{1 + \exp(-\varepsilon)}. \tag{19}$$

We solve the maximization of the log-likelihood function $LL(\boldsymbol{\sigma}, b, p, \boldsymbol{F})$ in three steps. In *Step 1*, we solve for the team's decision recursively for a given $(b, p, \boldsymbol{F})$, as described in the previous section. For the terminal version, we assume $\overline{T} = 25$. If, by this last period, a breakthrough has occurred and $\sigma_{\overline{T}} = 1$, then the project is successful and published in $t = 25 = \overline{T}$. If a breakthrough has not been observed and $\sigma_{\overline{T}} = 0$, then the project has failed and is abandoned. In this first step, we obtain critical values $\bar{\varepsilon}_t^{\sigma_t}$ for all $t = 1, \dots 24$ for the pre-breakthrough phase (with $\sigma_t = 0$) and the post-breakthrough phase (with $\sigma_t = 1$). In $t = 25 = \overline{T}$ no decision is taken. In the post-breakthrough phase, the critical value $\bar{\varepsilon}_t^1$ is such that $V(t+1) - F(t) - \bar{\varepsilon}_t^1 = \hat{\pi}(t)$ and $V(t) = \hat{\pi}(t)$ in equation (12). In the pre-breakthrough phase, this critical value $\bar{\varepsilon}_t^0$ is such that $E(V(t+1)|0) - F(t) - \bar{\varepsilon}_t^0 = 0$ and $V(t) = 0$ in equation (13). Through the recursive characterization of the team's optimization problem, we obtain a sequence $\{\bar{\varepsilon}_t^{\sigma_t}\}_{t=1}^{\overline{T}-1}$ and, using the CDF in equation (19), a sequence $\{G^{\sigma_t}(t)\}_{t=1}^{\overline{T}-1}$ with

$$G^{\sigma_t}(t) = G(\bar{\varepsilon}_t^{\sigma_t}). \tag{20}$$

Note that $G^0(0) = 1$ and $G^1(0)$ is not defined. These $G^{\sigma_t}(t)$ are the continuation probabilities in a given t with status σ_t.

In *Step 2*, we use the sequence of continuation probabilities, $\{G^{\sigma_t}(t)\}_{t=1}^{\overline{T}-1}$

with $G^0(0) = 1$, to calculate the log-likelihood $LL(\boldsymbol{\sigma}, b, p, \boldsymbol{F})$ in equation (18) for a vector of parameters $(b, p, \boldsymbol{F})$. *Step 3* runs optimization routines to find the vector $(b^*, p^*, \boldsymbol{F^*})$ with parameters that maximize the log-likelihood function,

$$(b^*, p^*, \boldsymbol{F^*}) \in \arg \max_{(b,p,\boldsymbol{F})\in\Omega} LL(\boldsymbol{\sigma}, b, p, \boldsymbol{F}) \tag{21}$$

with $\Omega = [0, 1]^2 \times \mathbb{R}^3$ and $\boldsymbol{F}$ such that $F(t) \geq 0$ for all $t = 1, \ldots, 24$.

5.2 Identification

Our model is based on a mixture of two optimal stopping problems, corresponding to the pre and post-breakthrough status of a given project. The breakthrough itself is an unobserved (to the econometrician) state variable until the time of publication. At that point, we can infer that a breakthrough occurred for each project that gets published as an RFC. However, we cannot infer that all abandoned projects are "bad" – some simply receive a large unobserved cost shock and exit before achieving a breakthrough. Intuitively, for each stopping problem, the share of projects that exit, either through publication or abandonment, identify whether the *net benefits* of continuation are positive. After a breakthrough, those net benefits consist of an observable marginal benefit $\hat{\pi}(t + 1) - \hat{\pi}(t)$ that is identified by the relationship between t and expected citations, as well as the "option value" associated with further improvements, less a marginal cost. Before a breakthrough occurs, all of the net benefits come in the form of option value, since the only reason to continue is in the hope of experiencing a breakthrough that would lead to publication and payoffs. The Bayesian learning process causes projects on the pre-breakthrough path to become more pessimistic about their option value over time, as the posterior belief that they are "bad" project increases.

In our preferred specification of the model, we assume that there is no learning ($b = 1$) and certain publication ($p = 1$) for nonstandards-track projects. 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. 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.[26] The second use of nonstandards is to provide information that complements a standard, such as guidelines for implementation and deployment.[27] Nonstandards-track RFCs require very little community agreement given their purely informational role.

We identify the payoffs in our model by estimating expected citations as a function of t. Although this could be done non-parametrically, in practice, we estimate a log-linear function, and let the payoffs vary for standards and nonstandards. Because there is no learning and no abandonment for nonstan-

[26]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.

[27]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).

dards, those projects can be used to identify all of the cost parameters in our empirical model by choosing an $F(t)$ such that the share of nonstandards published after t equals $1 - G(V(t+1) - \hat{\pi}(t) - F)$, i.e. the actual and implied probability of stopping are equal.

Given estimates of the payoffs and costs of continuation, the two parameters associated with the Bayesian learning process, p and b are identified by the rates of publication and abandonment. Intuitively, they are chosen to make the implied hazard rates line up with the empirical hazards depicted in the bottom left panel of Figure 1.

6 Results

We present our results for the estimated parameters b^*, p^*, and $\boldsymbol{F}^*$ that maximize the log-likelihood $LL(\boldsymbol{\sigma}, b, p, \boldsymbol{F})$. We fist consider our baseline model, assuming that all projects are ex ante identical. We then extend our preferred specification of the model by introducing project heterogeneity and estimate multiple values for b, p, and $\boldsymbol{F}$ for different project categories.

6.1 Baseline

Table 4 presents the estimated parameters for four different versions of our baseline model. We estimate our model on the full sample and the WG sample. For each of these samples, we present results for the estimation using the identification strategy based on standards-track and nonstandards-track RFCs splits as well as results without this identification strategy. We report the former in columns "Tracks" and the latter in columns "No Tracks".[28]

The full sample estimates are consistent with the presence of a majority of good projects, and a relatively small probability of breakthrough. These results suggest that IETF members enter the process believing that consensus is possible, as witnessed by an estimated value of good projects, p, hovering above $1/2$. However, they also expect that it will take them a relatively long time to achieve consensus, as reflected by an estimated value of the rate of learning, b, of about $1/4$.

6.1.1 Rate of Learning

In the pre-breakthrough phase in the full sample, one out of five good projects experiences a breakthrough in any given t, whereas it is one out of four for the WG sample. The rate of learning is higher for projects that are initiated within working groups relative to projects in the full sample. More specifically, projects sanctioned by working groups observe a breakthrough faster than individual projects. One possible explanation for this results is that the additional attention and feedback from the working group results in a higher rate of learning. We will provide some support for this in the section when we consider project heterogeneity.

6.1.2 Prior of Project Type

In the full sample, a bit less than 50% of projects are good and can, if a breakthrough occurs, generate value. The fraction of good projects is higher

[28]The reported parameters in column "Tracks" version are for standard-track projects.

Table 4: Baseline Results for Structural Model

	Full Sample		Working Group Sample	
	No Tracks	Tracks	No Tracks	Tracks
Learning (b)	0.27	0.18	0.37	0.28
	(0.004)	(0.004)	(0.007)	(0.007)
Priors (p)	0.56	0.45	0.68	0.55
	(0.003)	(0.004)	(0.004)	(0.005)
Costs $F(1)$	2.83	2.09	3.47	2.47
	(0.025)	(0.022)	(0.045)	(0.038)
Costs $F(10)$	0.85	0.86	0.81	0.89
	(0.005)	(0.006)	(0.007)	(0.008)
Costs $F(25)$	0.54	0.57	0.95	0.72
	(0.025)	(0.026)	(0.061)	(0.047)
Projects	16,268		3982	
(on/off standard track)		14,549/1719		3201/781
Versions	59,713		22,580	
(on/off standard track)		48,009/11,704		17,351/5229
Log-Likelihood	-39,128.6	-33,817.6	-12,521.6	-11,086.6
AIC	78,267.1	67,645.3	25,053.2	22,183.2

Standard errors in parentheses. “Tracks” indicate results using the standard-nonstandards-track approach with $b = 1 = p$ for nonstandards track projects; “No Tracks” estimates all five parameters for projects both on and off the standard track. $F(1)$, $F(10)$, and $F(25)$ are the estimated non-stochastic costs in $t = 1$, $t = 10$, and $t = 25$.

in the WG sample. A likely explanation for these differences is self-selection: while any individual can start a project outside a working group, the project threshold value for the WG to facilitate a project is higher. Also, working groups are formed to solve identified problems, and working group projects are more likely than individual or outside projects to relate to these problems, thus receiving more attention and ultimately support.

6.1.3 Costs of a Revision

Table 4 presents the per-period costs for $t = 1$, $t = 10$, and the final period $t = 25 = \overline{T}$.[29] Costs are strictly positive, decreasing and convex for both the full sample and the WG sample, as shown in Figure 4 that plots $F(t)$ for $t = 1, \ldots, 25$. The monotonic decrease in incremental costs may capture learning-by-doing effects. Alternatively, this cost pattern could be linked to the decreasing rate of textual change, as depicted in Figure 1. Smaller textual changes for later versions come at lower cost. Of course, the reverse is possible, too. Because less effort or time is spent on later versions (i.e., lower costs), later versions exhibit smaller textual changes.

6.1.4 Payoffs, Beliefs and Hazard Rates

The upper-left panel in Figure 5 illustrates the empirical pattern of posterior beliefs in the pre-breakthrough phase. Moreover, the upper-right panel presents the expected value of a project conditional on t, $E(V(t+1)|\sigma_t)$, in the two phases. Finally, in the lower-left panel in Figure 5, we plot the critical values $\bar{\varepsilon}_t^{\sigma_t}$ (as defined in (4)) for the pre-breakthrough phase ($\sigma_t = 0$) and the post-breakthrough phase ($\sigma_t = 1$), with the respective continuation probabilities $G^{\sigma_t}(t)$ (as defined in (5)) in the lower-right panel.

In the pre-breakthrough phase, both $E(V(t+1)|0)$ and the critical values $\bar{\varepsilon}_t^0$ follow a non-monotonic pattern. The non-monotonic values of $E(V(t+1)|0)$ are the result of a combination of three forces. First, the team knows that,

[29]Note, final period costs $F(25)$ never materialize because there is, by assumption, no decision in $t = 25 = \overline{T}$.

Figure 4: Cost Estimates (Baseline Models)

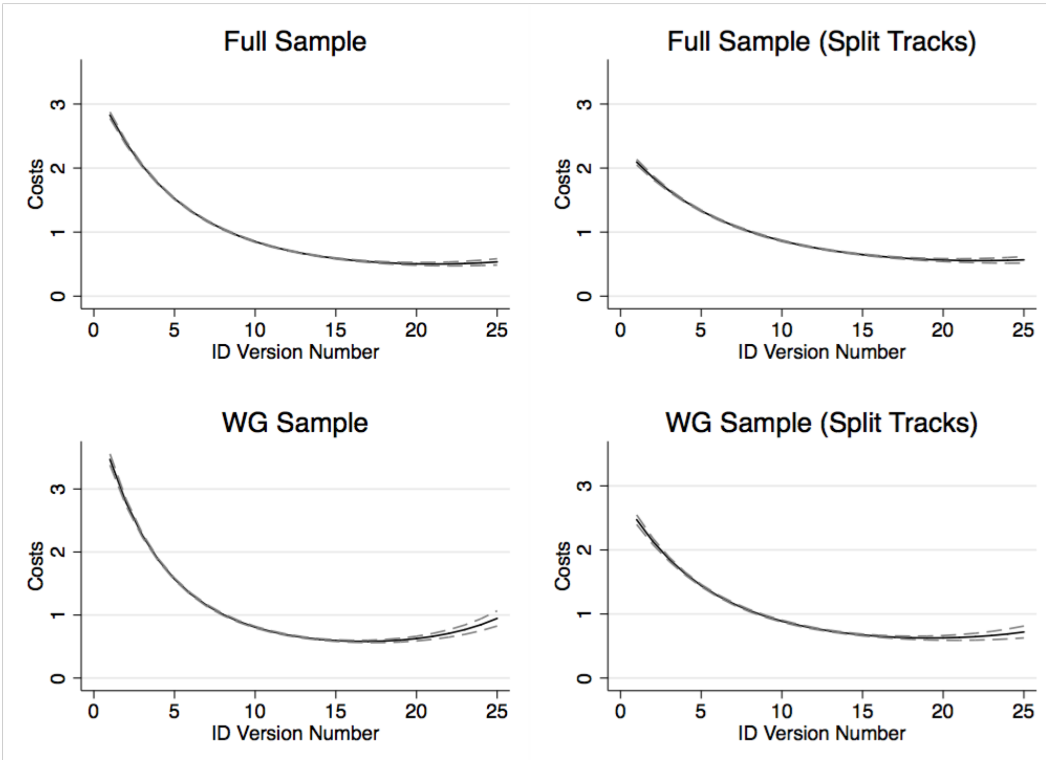

The four panels plot the cost estimates for the four versions of our baseline model in Table 4. Dotted lines are 95%-confidence bounds.

conditional on the realization of the breakthrough, the profits from publication increase in the number of versions. Second, the team anticipates that it will be paying lower values of the non-stochastic costs as the number of versions increases. Finally, the publication value is discounted by a value of the posterior beliefs that decreases with the version number.

Thus, as long as the increase in the expected value from publication, together with the reduction in the non-stochastic costs, compensate for the decrease in the posterior beliefs, the expected payoffs from continue increase with the version number. The expected values in $E(V(t+1)|0)$ decrease once the team becomes sufficiently pessimistic about the probability of the breakthrough and relative cost-savings, associated with higher versions, diminish. The same intuition explains the non-monotonic pattern of the critical values $\bar{\varepsilon}_t^0$ and the corresponding continuation probabilities, as depicted in the bottom panels of Figure 5.

In the post-breakthrough phase, beliefs do not play any role (as the team updates its posterior to $\hat{p}(t|1) = 1$ after the breakthrough). In this phase, ex-ante payoffs increase, tracking the increasing values of the publication payoffs, $\hat{\pi}(t)$, and the decrease in the value of non-stochastic costs, $F(t)$. The critical values $\bar{\varepsilon}_t^1$ monotonically decrease, which is explained by the fall in the option value of continuing. This follows the decreasing and convex pattern of the fixed costs, and the increasing and concave values of the payoffs $\hat{\pi}(t)$.

Finally, in Figure 6 we plot the conditional probabilities of outcomes. The panels on the left depict the hazards of the IETF data (full sample and WG sample); the panels on the right depict the hazards of simulated data using the estimated parameters in Table 4.[30] Our simulated hazard rates track both pattern and magnitude for the full sample. For the WG sample, the hazard rates for the simulated data track well the hazard rates for the IETF data. However, our estimates are not able to match the magnitude of the IETF data hazard rates.

[30]We simulate 20,000 on-track projects.

Figure 5: Decisions (Baseline Models)

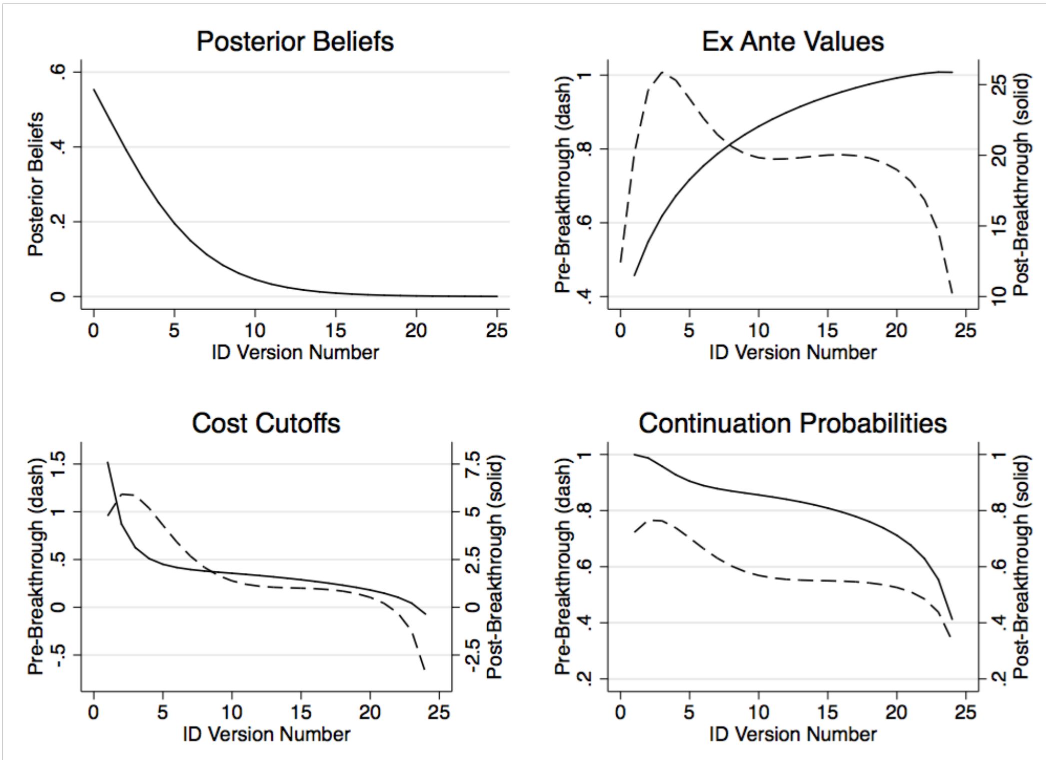

Estimated parameters for the baseline model in Table 4. Top-left: the team's posterior beliefs during the pre-breakthrough phase. Top-right: continuation values (pre-breakthrough and post-breakthrough) over ID Version Number. Bottom-left: the cutoff values for cost shocks, $\bar{\varepsilon}_t^{\sigma_t}$. Bottom-right: the continuation probabilities in $t.1$, before the team observes its cost shock ε_t.

6.2 Heterogeneity

For the baseline results, we have assumed that all IETF projects are ex ante identical. For the results below, we introduce ex-ante project heterogeneity and estimate different values for our model parameters. We focus on the WG sample with standard-track splits ("Tracks"). We summarize the results in Table 5. The first column reproduces the results from Table 4.

Figure 6: Hazard Rates (Data and Simulations)

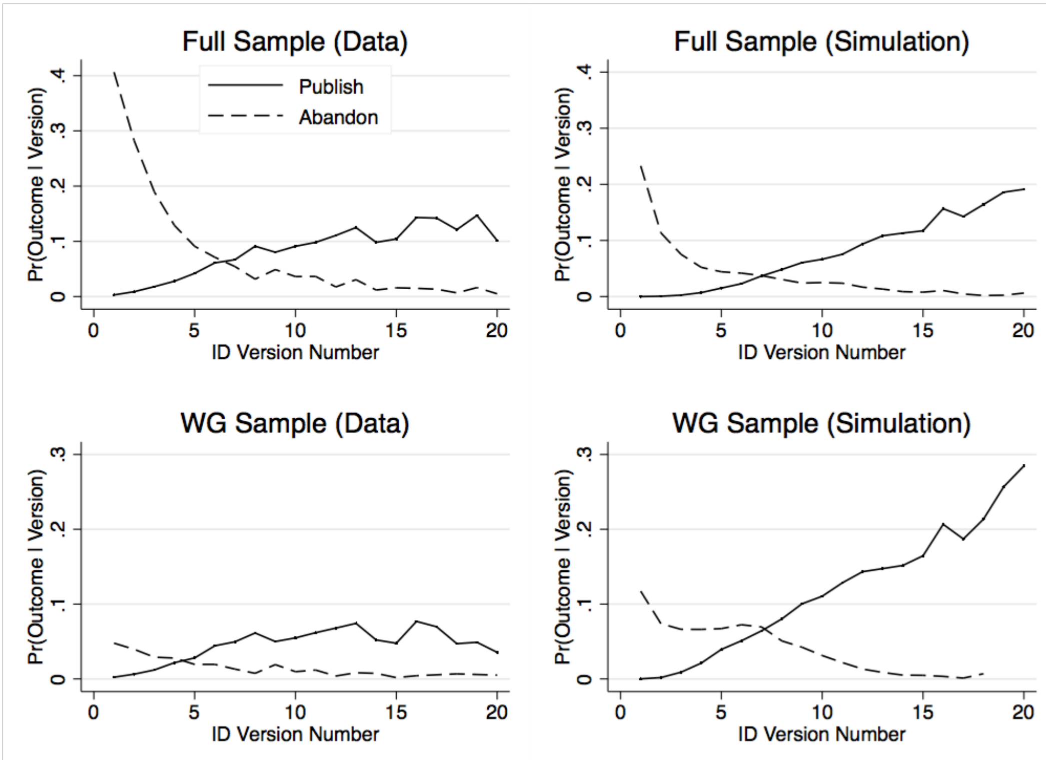

The four panels plot the hazard rates for published (solid) and abandoned (dashed) projects for the four versions of our baseline model in Table 4.

6.2.1 Rate of Learning

For the first model extension, we ask how project-related communication drives learning. We let b vary with the amount of attention and feedback a project receives. We hypothesize that more attention (via more project-related communication) is associated with a higher learning rate. We measure communication (or attention) using the number of e-mail messages per version sent during the revision process. Each project is assigned the mean of e-mail messages per version for each of the four quartiles. This gives us four categories: low, low-high, high-low, and high. For this exercise, we assume that neither the

Table 5: Heterogeneity Results for Structural Model

	Baseline	Emails (b)	Years (p)	Authors (F)
Learning (b)	0.278		0.278	0.315
	(0.007)		(0.008)	(0.007)
Priors (p)	0.553	0.560		0.630
	(0.005)	(0.005)		(0.005)
Costs, $F(1)$	2.473	2.323	2.450	
	(0.038)	(0.038)	(0.038)	
Costs, $F(10)$	0.888	0.913	0.891	
	(0.008)	(0.008)	(0.008)	
Costs, $F(25)$	0.719	0.617	0.706	
	(0.047)	(0.041)	(0.047)	
Learning (b): low		0.189		
		(0.007)		
Learning (b): low-high		0.209		
		(0.007)		
Learning (b): high-low		0.246		
		(0.007)		
Learning (b): high		0.365		
		(0.015)		
Prior (p): 1996–2000			0.520	
			(0.006)	
Prior (p): 2001–2005			0.570	
			(0.006)	
Prior (p): 2006–2009			0.616	
			(0.008)	
Costs, $F(1)$: 1 author				3.156
				(0.054)
Costs, $F(1)$: 2 authors				3.118
				(0.054)
Costs, $F(1)$: 3–4 authors				3.023
				(0.054)
Costs, $F(1)$: 5+ authors				2.741
				(0.052)
Projects (on/off standard track)	3201/781			
Versions (on/off standard track)	17,351/5229			
Log-Likelihood	-11,086.6	-11,016.5	-11,011.4	-9,873.0
AIC	22,183.2	22,045.0	22,036.8	19,762.1

Standard errors in parentheses. Estimates for WG sample with standard-track splits.

prior p nor the non-stochastic cost component $\boldsymbol{F}$ depend on communication.

We find that more attention increases the estimated value for b; it thus increases the rate of learning. This in return implies that more attention induces faster updating of beliefs. Players thus become pessimistic more rapidly. At the same time, a breakthrough, if it does occur, arrives faster which results in higher continuation value in the pre-breakthrough phase. In the first row in Figure 7, we plot the continuation probabilities $G(\bar{\varepsilon}_t^{\sigma_t})$ for three estimated values of b (low, high-low, and high) to capture this compound effect. A lower probability of continuation implies faster or earlier stopping. We can see that a higher value of b does not affect the continuation probabilities in the post-breakthrough phase. In the pre-breakthrough phase, except for the first few versions, a faster rate of learning induces faster rate of quitting.

6.2.2 Prior of Project Type

In Table 2, we see a varying number of projects per year for different periods. Moreover, projects initiated in earlier years receive more patent citations per year than younger projects. This suggests circumstances at the IETF that change over time. We consider three different periods (1996–2000, 2001–2005, and 2006–2009) and estimate the prior probability p for each of these periods. For this exercise, we assume that neither b nor $\boldsymbol{F}$ change over time.

The results in Table 5 illustrate that, over the years, the prior probability that a project is good has increased. There are at least two competing explanations for this. First, projects may have become inherently better, lifting the prior. Alternatively, the IETF may have become more lenient, endorsing

Figure 7: Continuation Probabilities

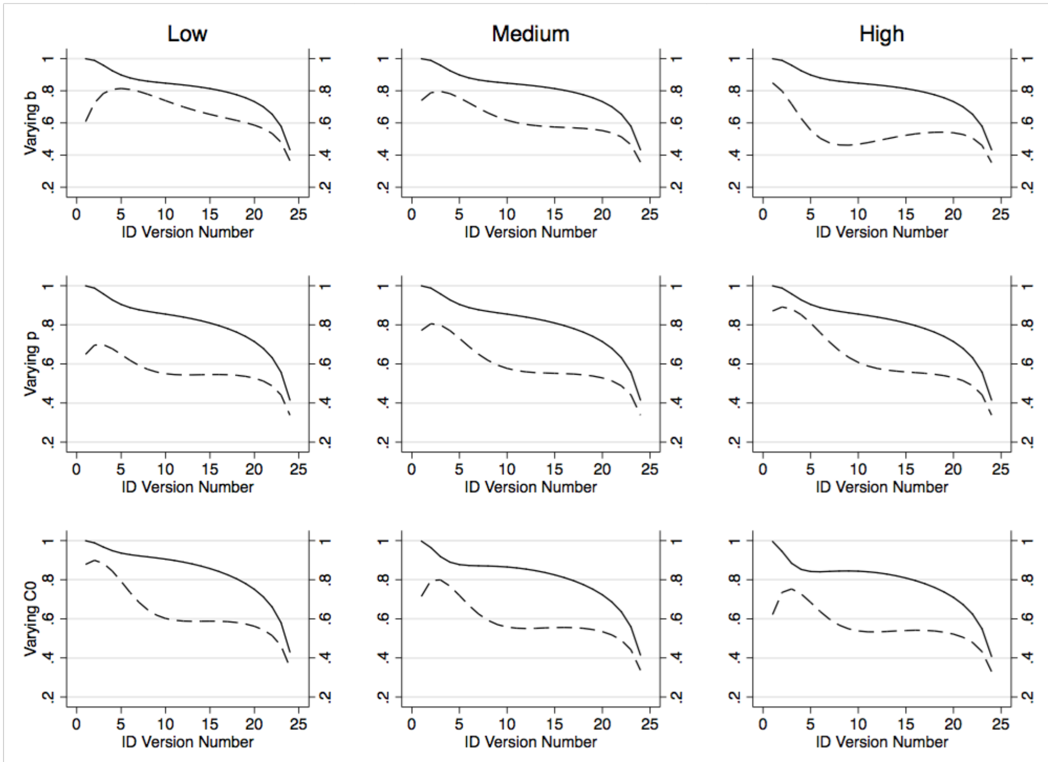

We plot continuation probabilities for the pre-breakthrough phase (dashed) and the post-breakthrough phase (solid) for three different values of b (first row: low, high-low, and high), p (second row: 1996–2000, 2001–2005, and 2006–2009), and C_0 (third row: 1 author, 3–4 authors, and 5+ authors). Estimated parameter values in Table 5.

more projects as RFCs. This in return increases the prior p because only good projects can be published. Note that a higher value of p induces less pessimistic teams in the pre-breakthrough phase. In addition, a higher value of p implies higher continuation values. In the second row of Figure 7, we see the compound effect. Again, in the post-breakthrough phase, the continuation probabilities are not driven by p, because posteriors are $\hat{p}(t|1) = 1$ for all t. In the pre-breakthrough phase, however, continuation probabilities are higher for higher values of p. As teams expect projects to be better, they continue to submit new versions and quit later.

6.2.3 Costs of a Revision

We see in Table 2 that projects by teams with more authors receive more patent citations. This is partly because larger teams have longer projects. Moreover, the graphs in Figure 1 suggest that projects with more versions receive more patent citations. One possible explanation for the positive relationship between team size and version is that larger teams have lower costs. We find weak support for this. The costs of a first revision are lower when the team has more authors. In Figure 8, we plot the costs against ID Version Number; for later versions there is no statistical difference; author team size matters for costs only early in the revision process. Albeit small, the differences in the cost-intercept C_0 do matter for the team's continuation probabilities. The third row of Figure 7 plots the continuation probabilties for three different cost intercepts. As the non-stochastic costs of another version increase, the continuation probabilities (before teams observe their costs shocks) decrease.

7 Counterfactual Analysis

For our first counterfactual analysis, we vary the IETF's quality standards (with respect to endorsed projects) by varying the value for p. We use the estimated parameters for the WG sample from our preferred model (using the standards-nonstandards track approach) as presented in the last column in Table 4. We consider two different approaches. First, we vary p while keeping

Figure 8: Cost Estimates (Heterogeneity Results)

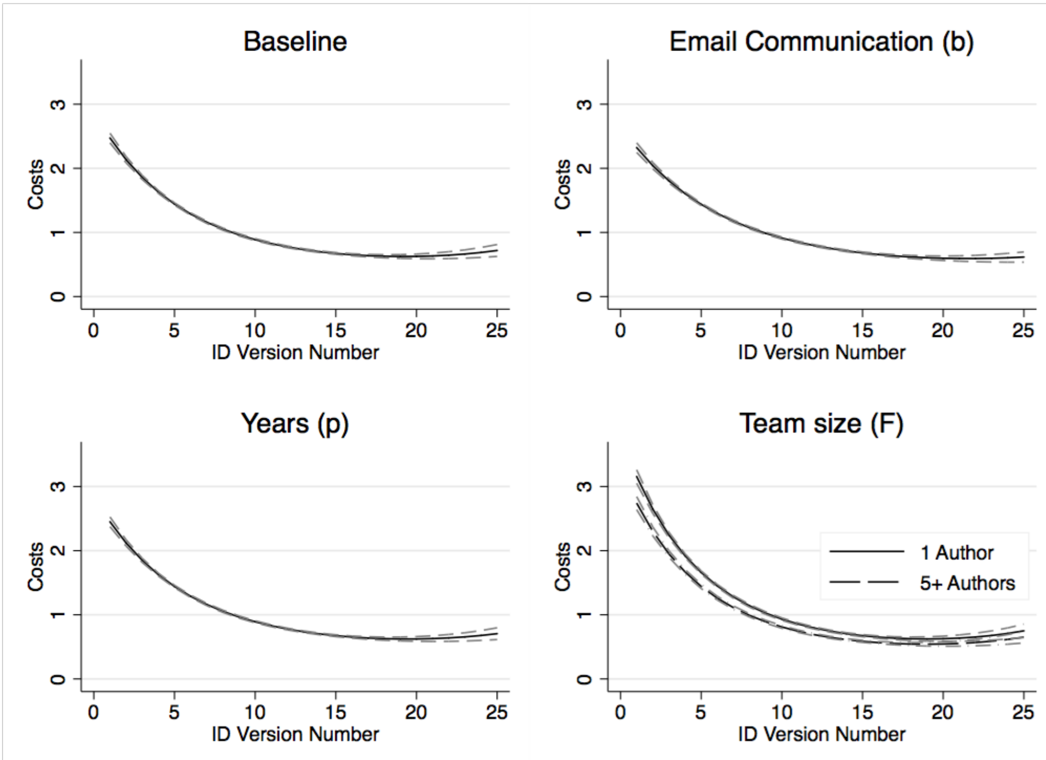

The four panels plot the cost estimates for our baseline model and three heterogeneity extensions in Table 5. In the lower-right panel, we plot the costs for two different team sizes. Dotted lines are 95%-confidence bounds.

the number of projects constant. For the results reported below, we simulate data for 20,000 projects. This approach allows for both varying number of overall versions and of the total costs incurred by the author teams. For the second approach, we vary p while keeping the total realized costs of all projects constant.[31] This implies a varying number of projects and versions but accounts for possible cost budget. Note that we do not keep the number of versions constant, because, as depicted in Figures 4 and 8, versions come at

[31]The total realized costs of 20,000 projects, given the estimated parameters, is 95,619. For this second approach, we round up and assume a total cost budget of 100,000. The actual number of projects is such that total costs do not exceed the cost budget.

Figure 9: Counterfactuals (Prior p)

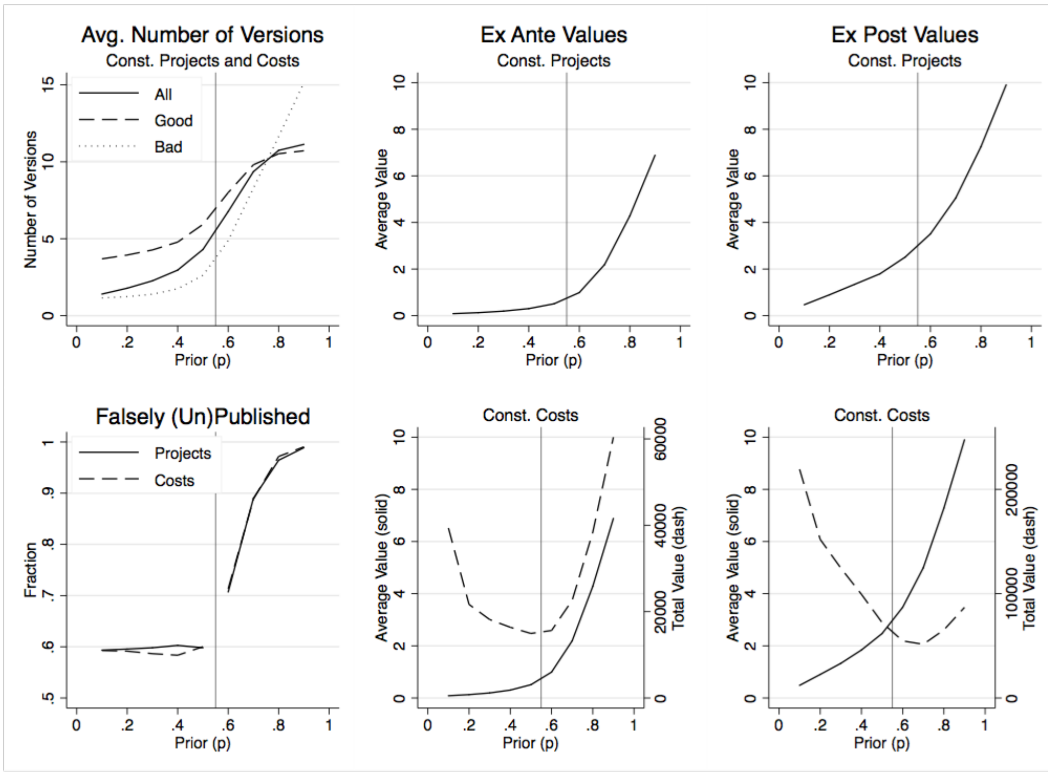

Panels depict results from the first counterfactual exercise. We vary p from $^1/_{10}$ to $^9/_{10}$.

varying cost where earlier versions are more expensive than later version.

In Figure 9, we illustrate the results from our first set of counterfactuals. Varying the prior probability p reflects varyings degree of leniency by the IETF. In the top-left panel, we plot the average number of versions per project separately for all (solid), good (dashed), and bad (projects). As the prior p increases and the IETF becomes more lenient, accepting more projects as good projects, the average number increases. Observe that good projects undergo more versions than bad projects when p is low but fewer version when p is high. Bad projects do not experience a breakthrough; but with high values of p, the continuation value is high relative to costs (which are relatively low at high version numbers). This is because of the discontinuity in values at a breakthrough—implying high costs of stopping. Good projects, after a breakthrough, have a lower continuation value because patent citations are concave in version number (and flatten out). Teams continue to submit new versions longer for projects that have not yet experienced a breakthrough.

The bottom-left panel plots falsely published and unpublished projects (for both constant project numbers and costs). Recall, we treat the estimated parameter p^* as the true value. If, as in our counterfactuals, the IETF chooses a value $p < p^*$ [$p > p^*$] it is stricter [more lenient] than under this true value. All projects $p' \in [p, p^*)$ [$p' \in (p^*, p]$] are *treated* projects that are falsely considered bad [good]. Falsely bad projects are not published when they would otherwise potentially be published under $p = p^*$. For all $p < p^*$, we plot the thus unpublished projects as a share of the falsely bad projects. Falsely good projects may be published when they would otherwise not be published under $p = p^*$. For $p > p^*$, we plot the thus published projects as a share of the falsely good projects. We can see that for treated projects, the fraction of falsely (un)published projects is constant when p is low but increasing when p is high. In fact, for $p = 0.9$ almost all treated projects (falsely good projects) are published. This means that the errors stemming from too strict an IETF (for $p < p^*$) are mitigated by shorter processes so that not all treated projects are published. The same errors stemming from too lenient an IETF (for $p > p^*$) are all materialized as a projects are longer and the error rate (unpublished good projects) goes to zero.

In the center column of Figure 9, we plot ex-ante values (average and total) for constant project numbers (top) and constant costs (bottom). For both constant project numbers and constant costs, the average ex-ante value (a team's continuation value in $t = 0$, before the first version is submitted) is increasing in p. For the constant cost numbers, the total ex-ante value is U-shaped. This is a result of the assumption that the number of projects is chosen to keep the total costs constant. For low values of p with few version, the number of projects is high (with low ex-ante value per project) where the number of projects is low (with high ex-ante value per project) otherwise. The final result is a picture that suggests (as institutional choice) either a relatively strict or relatively lenient IETF, with the estimated parameter of p^* generating low total ex-ante values.

In the right column of Figure 9, we plot ex-post values (average and total) for constant project numbers (top) and constant costs (bottom). The picture for constant project numbers is analogous to the one for ex-ante values. We also find a U-shaped relationship between total ex-post costs and the prior p. Unlike for ex-ante values, however, lower values of p now dominate higher values of p.

8 Concluding Remarks

We propose a model of research and development as a process of experimentation in which researchers repeatedly revise specifications of a project and update their beliefs about the project's type. Only a good project whose type is learned by researchers can generate value. Researchers abandon a project when the opportunity costs of continuing exceed the expected benefits. We estimate the structural parameters of this dynamic optimization problem using a novel data set with information on both successful and abandoned projects from the Internet Engineering Task Force (IETF), a standard development organization that creates and maintains standards necessary for the functioning of the internet. The structural approach allows us to recover the researchers' unobserved beliefs and opportunity costs, and answer questions about whether specific rules and institutions encourage "efficient abandonment" by researchers. We find that opportunity costs are decreasing over time and feedback and comments from the IETF community at large increase the speed at which developers learn the type and potential of a project. A higher rate of learning reduces the costs of extensive and fruitless development of bad projects (without value).

References

Allen, T. J. (1966): "Studies of hte Problem-Solving Process in Engineering Design," *IEEE Transactions on Engineering Management*, 13, 72–83. (Not cited.)

Bergemann, D. and J. Välimäki (2008): "Bandit Problems," in *The New Palgrave Dictionary of Economics, Second Edition*, ed. by S. N. Durlauf and L. E. Blume, Palgrave Macmillan. (Not cited.)

Bloom, N., R. Sadun, and J. Van Reenen (2012): "Americans Do IT Better: US Multinationals and the Productivity Miracle," *American Economic Review*, 102, 167–201. (Not cited.)

Bloom, N., M. Schankerman, and J. Van Reenen (2013): "Identifying Technological Spillovers and Product Market Rivalry," *Econometrica*, 81, 1347–1393. (Not cited.)

Crawford, G. and M. Shum (2005): "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 73, 1137–1173. (Not cited.)

Ganglmair, B. and E. Tarantino (2014): "Conversation with Secrets," *RAND Journal of Economics*, 45, 273–302. (Not cited.)

Heidhues, P., S. Rady, and P. Strack (2015): "Strategic Experimentation with Private Payoffs," *Journal of Economic Theory*, 159A, 531–551. (Not cited.)

Hellmann, T. and E. Perotti (2011): "The Circulation of Ideas in Firms and Markets," *Management Science*, 57, 1813–1826. (Not cited.)

Keller, G., S. Rady, and M. Cripps (2005): "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73, 39–68. (Not cited.)

Lerner, J. (1999): "The Government as Venture Capitalist: The Long-Run Effects of the SBIR Program," *Journal of Business*, 72, 285–318. (Not cited.)

Lerner, J. and S. Stern (2010): *Innovation Policy and the Economy: Volume 10*, National Bureau of Economic Research, Inc. (Not cited.)

Pakes, A. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 55, 911–922. (Not cited.)

Rothschild, M. (1974): "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9, 185–202. (Not cited.)

Rysman, M. and T. Simcoe (2008): "Patents and the Performance of Voluntary Standard Setting Organizations," *Management Science*, 54, 1920–1934. (Not cited.)

Simcoe, T. (2012): "Standard Setting Committees: Consensus Governance for Shared Technology Platforms," *American Economic Review*, 102, 305–336. (Not cited.)

Stein, J. (2008): "Conversations among Competitors," *American Economic Review*, 98, 2150–2162. (Not cited.)

Wallsten, S. (2000): "The Effects of Government-Industry R&D Programs on Private R&D: The Case of the Small Business Innovation Research Program," *RAND Journal of Economics*, 4, 82–100. (Not cited.)

WUCHTY, S., B. F. JONES, AND B. UZZI (2007): "The Increasing Dominance of Teams in Production of Knowledge," *Science*, 316, 1036–39. (Not cited.)