Learning When to Quit: An Empirical Model of Experimentation

Bernhard Ganglmair   Timothy Simcoe   Emanuele Tarantino

January 2018

This Paper

Learning in R&D

How do scientists/engineers learn whether or not they have a good idea?

- Study dynamic learning in R&D decision making at the project level
- Learning about the quality (or type: good or bad) of a project that materializes via "consensus" by a community of peers
- Parsimonious structural model that combines Bayesian learning with ongoing improvements to project
- Estimate on project-level data for successful and abandoned projects from the *Internet Engineering Task Force* (IETF)

How do Researchers and Engineers Use Time?

Ideally, focus on valuable projects and quickly cut loose projects without potential. But how to know which is which? (Allen, 1966)

- Researchers' actions (i.e., employment of time) depend on beliefs and the (institutional) environment
- This paper:
  1. Estimate key parameters in researchers' R&D decisions (beliefs, opportunity costs) based on observed actions and outcomes
  2. Counterfactual: Are cost subsidies or publication prizes the better policy?
  3. Counterfactual: What is the effect/cost of over-confidence?
  4. Heterogeneity: How do communication and prior experience relate to learning?

Summary of Findings

1. Estimates for "rate of learning," prior beliefs about quality, and opportunity costs of revisions
   - Three out of five (59%) projects submitted to IETF are publishable (can generate value)
   - Less than 1/3 (31%) of these projects are published (successful)
   - Realization of value for 17% of good projects per revision
   - Decreasing opportunity costs of revising a project's draft

2. Both *subsidies* and *prizes* raise projects' *gross* value, only the latter increase *net* value (i.e., considering private revision costs)

3. Over-confidence is more costly than pessimism (misaligned priors)

4. More communication and experience of authors result in faster learning; higher commerciality of projects slows the rate of learning

Related Literature

- Dynamic decisions with learning
    - Erdem and Keane (1996); Crawford and Shum (2005); Dickstein (2014)
    - Pakes (1986); Nanda and Rhodes-Kropf (2015); Krieger (2017)
- Consensus standardization (IETF)
    - Rysman and Simcoe (2008); Fleming and Waguespack (2009); Simcoe (2012); Ganglmair and Tarantino (2014)
- Prizes v. Subsidies
    - Murray, Stern, Campbell, and MacKormack (2012); Galasso, Mitchell, and Virag (2016); Hall and Van Reenen (2000)
- Optimism in innovation
    - Galasso and Simcoe (2011); Hirshleifer, Low, and Hong Teoh (2012)

Institutions

Internet Engineering Task Force

Institutional Background

- IETF is the main forum for internet protocol development
- Anyone can participate. In practice, corporate, academic and individual engineers, and computer scientists
- Main motivation: Advance technology
- Transparent process $\Rightarrow$ rich data
  - Repository with every version of every project (success and failure)
  - E-mail server where project-related communication occurs
  - Tri-annual meetings, held around the world

IETF Protocol Examples

Important Standards

	Description	Year
RTP	Real-time Transport Protocol	2003
SIP	Session Initiation Protocol	2002
HTTP	Hypertext Transfer Protocol	1999
IPV6	Internet Protocol, Version 6 (IPv6)	1998
DHCP	Dynamic Host Configuration Protocol	1997
POP3	Post Office Protocol – Version 3	1996
NAT	Network Address Translator	1994
FTP	File Transfer Protocol	1985
TCP	Transmission Control Protocol	1981
IP	Internet Protocol	1981

▸ Major Contributors

IETF Standardization

Process Overview: Data-Generating Process

1. Identify problem and submit proposal (Internet Draft or ID)
   - Two types: Individual and Working Group
   - All projects posted to public repository

2. Community feedback via email and meetings

3. **Rough consensus** $\Rightarrow$ ID published
   - Decision by WG Chair and IESG (*de facto* super-majority)
   - Published ID's called Proposed Standards (or RFCs)

4. No consensus $\Rightarrow$ sponsors have a choice
   - Revise ID $\rightarrow$ return to step (1) [submit revision]
   - Abandon ID $\rightarrow$ expires in 6 months

Standards and Nonstandards

Parallel Publication Tracks

- Standards
  - Normative document ⇒ formal IETF endorsement
  - Objective: promote commercial implementation
  - Example: “Domain Host Control Protocol” ▸ RFC2131
- Nonstandards
  - Informative document
  - Two types: Informational & Experimental
  - Example: “Known TCP Implementation Problems” ▸ RFC2525
- Identical review and publication process
  - Nonstandards have less uncertainty → “no learning” baseline

IETF Standardization

Possible Outcomes

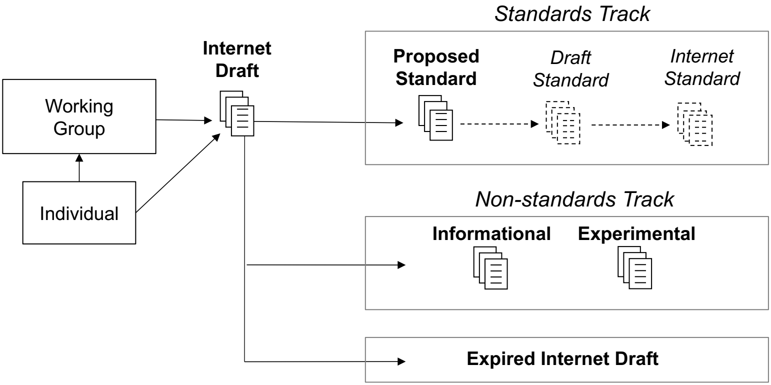

Data & Descriptive Analysis

```
HTTP Working Group                                   R. Fielding
INTERNET-DRAFT                                         UC Irvine
<draft-ietf-http-v11-spec-rev-00.txt>                  J. Gettys
                                                     J. C. Mogul
                                                             DEC
                                                      H. Frystyk
                                                  T. Berners-Lee
                                                         MIT/LCS
Expires January 30, 1998                           July 30, 1997
```

Hypertext Transfer Protocol -- HTTP/1.1

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or made obsolete by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress".

To learn the current status of any Internet-Draft, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Distribution of this document is unlimited. Please send comments to the HTTP working group at <http-wg@cuckoo.hpl.hp.com>. Discussions of the working group are archived at <URL:http://www.ics.uci.edu/pub/ietf/http/>. General discussions about HTTP and the applications which use HTTP should take place on the <www-talk@w3.org> mailing list.

Abstract

The Hypertext Transfer Protocol (HTTP) is an application-level protocol for distributed, collaborative, hypermedia information systems. It is a generic, stateless, object-oriented protocol which can be used for many tasks, such as name servers and distributed object management systems, through extension of its request methods. A feature of HTTP is the typing and negotiation of data representation, allowing systems to be built independently of the data being transferred.

Fielding, et al [Page 1]

▸ URL

HTTP Working Group — R. Fielding
INTERNET-DRAFT — UC Irvine
<draft-ietf-http-v11-spec-rev-00.txt> — J. Gettys
J. C. Mogul
DEC
H. Frystyk
T. Berners-Lee
MIT/LCS
Expires January 30, 1998 — July 30, 1997

Hypertext Transfer Protocol -- HTTP/1.1

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

URL

Hypertext Transfer Protocol -- HTTP/1.1

RFC 2616

Document | IESG evaluation record | IESG writeups | Email expansions | History

Revision differences

From revision draft-ietf-http-v11-spec-rev-06 (1998-11-24)

To revision rfc2616 (1999-06-01)

Diff format Side-by-side | Before-after | Change bars | Wdiff   Submit

Document history

Date	Rev.	By	Action
1999-06-01	RFC	(System)	RFC published
1999-02-25	06	(System)	IESG has approved the document
1998-11-24	06	(System)	New version available: draft-ietf-http-v11-spec-rev-06.txt (diff from previous)
1998-09-30	05	(System)	Last call sent
1998-09-15	05	(System)	New version available: draft-ietf-http-v11-spec-rev-05.txt (diff from previous)
1998-08-14	04	(System)	New version available: draft-ietf-http-v11-spec-rev-04.txt (diff from previous)
1998-04-07	03	(System)	New version available: draft-ietf-http-v11-spec-rev-03.txt (diff from previous)
1998-02-24	02	(System)	New version available: draft-ietf-http-v11-spec-rev-02.txt (diff from previous)
1997-11-26	01	(System)	New version available: draft-ietf-http-v11-spec-rev-01.txt (diff from previous)
1997-10-16	00	(System)	New version available: draft-ietf-http-v11-spec-rev-00.txt

▸ URL

Main Variables

- Decisions: **Event** in $t \in$ {Continue, Abandon, Publish}
- **Citations** $\Leftrightarrow$ E[π|Publish, Version]
  - Count of non-patent prior art references to RFC
  - *Alternative: Count of RFC references to previous RFC*
- Demographics
  - Cohort, team size (authors), experience (max{previous IDs})
- Feedback
  - Count of emails that specifically mention ID
- Revision size
  - Cosine distance for document-pair word vectors

Estimation Sample

IETF Submissions: 1996–2009

	Full Sample	Working Group	Stds-track Aband.	Stds-track Publ.	Nonstds-track
WG (%)	24.44	100.00	14.11	65.25	46.39
Team Size (Author Count)	2.28	2.45	2.22	2.43	2.48
Experience (max Projects)	15.01	15.69	13.50	21.87	16.84
Versions	3.55	5.60	2.09	9.33	6.71
Communication (Emails)	21.20	33.78	9.87	69.19	41.03
Published RFC (%)	23.97	56.10	0.00	100.00	100.00
Citations	2.99	8.30	0.76	12.23	7.19
N (Projects)	16,091	3,932	12,234	2,210	1,647
N (Versions)	57,179	22,025	25,511	20,622	11,046

- Author-teams on individual and WG submissions have similar size and experience; but strong correlation between experience and publication outcome
- WG proposals account for 25% of all IDs, but have higher publication rates
- Published IDs go through more revisions than abandoned IDs
- Patent citations increase with publication as RFCs; standards-track proposals are cited more

Mean Distance to Original Proposal

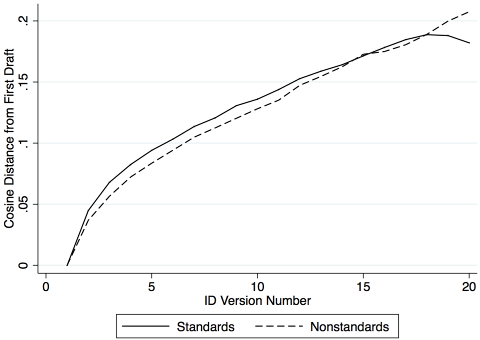

- Textual distance of a version T from initial version $t = 1$
- Proposals change throughout the revision process
- Plot very similar for standards-track and nonstandards-track proposals

$$\text{dist}(T, 1) \equiv 1 - \frac{x_T \cdot x_1}{||x_T|| \; ||x_1||} \qquad \text{with } ||x_i|| = \sqrt{\sum_{l=1}^{n} x_{i,l}^2}\,.$$

Hazard Rates

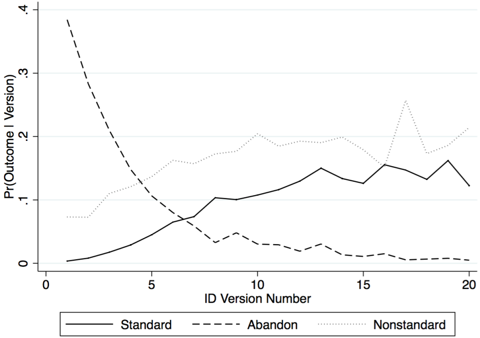

- Probability of exit, conditional on survival to t
- 40% of proposals never revised
- Some immediate publication of nonstandards (no learning)
- Publication hazard (for standards-track) always below 16%: Authors continue to revise ⇒ 2-phase model

Citations Increase with Revisions

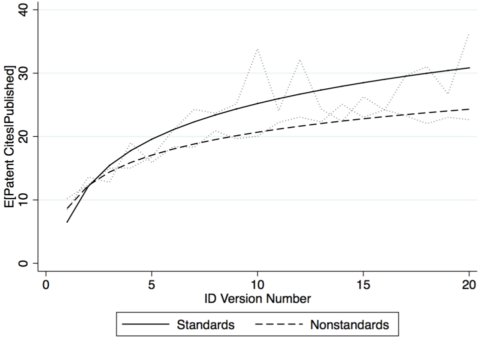

- Citations from U.S. patents (NPL) as proxy for commercial impact
- Citations increase with number of revisions
- Lower for nonstandards-track RFCs than standards

$$Cites_i = \alpha_y + (\beta_1 + \beta_2 Nonstandard_i) + \log(Versions_i) + \varepsilon_i$$

Implications

- Without learning, linear hazard rates of abandonment . . .
- . . . but data suggests process is not *purely* about learning
  - Distance measure → cumulative improvement
- Data suggest two phases: before and after consensus
  - Nonstandards also take time to publish: not all are immediately published
  - Decreasing hazard rates of abandonment, increasing hazard rates of publication
- Empirical goals:
  - Recover primitives: opportunity costs, learning parameter, prior belief of project quality
  - Replicate hazard rates

Model & Estimation

Bayesian Learning Model

- Model of R&D in the spirit of Roberts and Weitzman (1981)
- Researchers develop a project; endowed with an initial version/idea of unknown quality (good/bad)
- Only good projects generate value; good type must be learned (via breakthrough ~ consensus)
- Two-phase model:
  - Costly experiments to learn the project's type (→two-armed bandit)
  - Once type is learned in a period τ, further project revisions to increase value (→stopping problem)
- In each period, after observing a cost shock $\varepsilon_t \sim G(\cdot)$, the team decides to *continue* or *stop*

Bayesian Learning Model

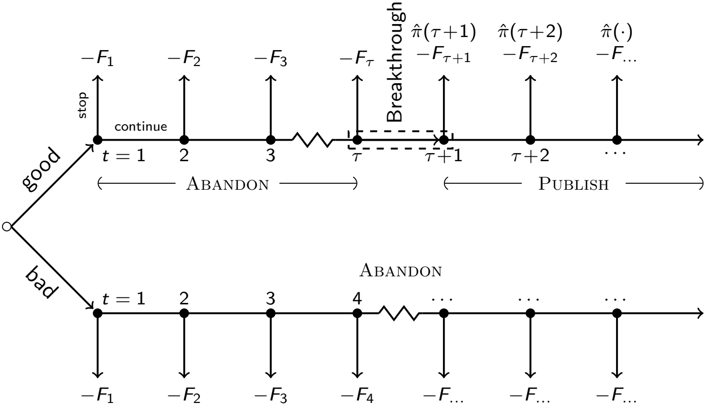

$$F_t = \sum_{k=0}^{t-1} F(k) + \varepsilon_k, \qquad F(0) = \varepsilon_0 = 0$$

Stylized Model: Two-Period Example

Illustrate; and Motivate Empirical Identification

- p: prior belief/probability that project is good
- b: probability that researcher learns good type (breakthrough) after posting a revision ($= 0$ if bad project)
- $\hat{\pi}_t$: payoffs when published in t; $\hat{\pi}_1 < \hat{\pi}_2$
- (A_t, S_t): share of abandoned/published in t: $A_1 + A_2 + S_1 + S_2 = 1$
- Two-period model timing:
  - $t = 0$: researcher endowed with initial proposal (good with p)
  - $t = 1$: researcher observes ε_1; decides to *continue* or *stop* (allowing for learning with b)
  - $t = 2$: if good type revealed, project is success (post-breakthrough), otherwise failure (pre-breakthrough)

good project with p: bad project with $1-p$:

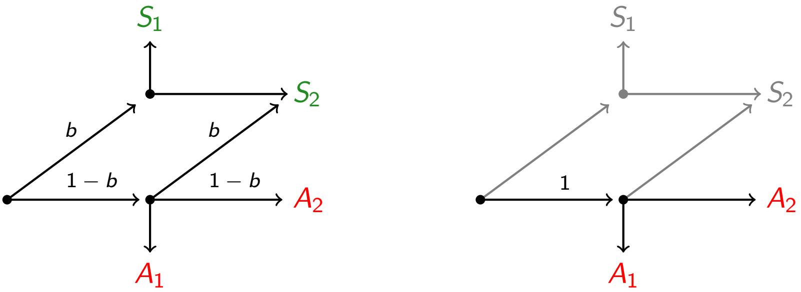

Beliefs in $t = 1$ are updated using Bayes' rule:

$$\hat{p}^{\text{pre}} = \frac{p(1-b)}{1-pb},$$

for pre-breakthrough, and $\hat{p}^{\text{post}} = 1$ for post-breakthrough.

Trade-Offs

- **Before breakthrough** (not yet learned good type):
  - Beliefs: Failed experiments ⇒ researchers become more pessimistic

$$\hat{p}^{\text{pre}} < p$$

  - Benefits: More revisions ⇒ larger potential publication payoff

$$\hat{\pi}_2 > \hat{\pi}_1 > 0$$

  - Tradeoff: *Exploitation vs. exploration*
  - Continue if and only if

$$b\hat{p}^{\text{pre}}\hat{\pi}_2 - [F^{\text{pre}}(1) + \varepsilon_1] \geq 0 \qquad \rightarrow \qquad G(b\hat{p}^{\text{pre}}\hat{\pi}_2 - F^{\text{pre}}(1))$$

Trade-Offs

- **After breakthrough** (learned good type):
  - Beliefs: Good type has been learned so that $\hat{p}^{\text{post}} = 1$
  - Tradeoff: Simple comparison of marginal costs and benefits
  - Continue if and only if

$$\hat{\pi}_2 - \left[F^{\text{post}}(1) + \varepsilon_1\right] \geq \hat{\pi}_1 \qquad \rightarrow \qquad G(\hat{\pi}_2 - \hat{\pi}_1 - F^{\text{post}}(1))$$

From Model to Identification

- Use conditional choice probabilities together with data on payoffs and distribution of the structural error terms to identify cost functions:

$$\Pr(\text{stop}|\text{post}) = \frac{S_1}{N_1^{\text{post}}} = 1 - G(\hat{\pi}_2 - \hat{\pi}_1 - F^{\text{post}}(1))$$

$$\Pr(\text{stop}|\text{pre}) = \frac{A_1}{N_1^{\text{pre}}} = 1 - G(b\hat{p}^{\text{pre}}\hat{\pi}_2 - F^{\text{pre}}(1))$$

where

$$N_1^{\text{post}} = S_1 + S_2 - A_2 \frac{b\hat{p}^{\text{pre}}}{1 - b\hat{p}^{\text{pre}}}$$

$$N_1^{\text{pre}} = A_1 + A_2 \frac{1}{1 - b\hat{p}^{\text{pre}}}$$

are the number of projects in the post- and pre-breakthrough phases.

From Model to Identification

Theorem

For any $b \geq \underline{b}$, there exists a unique $p(b)$ and $\boldsymbol{F}(b)$ such that for all t, the probability of publication (abandonment) in period t equals S_t (A_t).

- For any admissible b, we can choose other structural parameters to rationalize the data
- For point identification, we need more information
- We consider the cost functions!

From Model to Identification

Assumption

Revisions costs are independent of the breakthrough phase:

$$F^{post}(1) = F^{pre}(1) = F(1)$$

Then:

Theorem

If $G(\cdot)$ is the logistic cdf and given Assumption 1, there exists a unique solution so long as $p \geq 1 - A_1^2$.

- We establish this for a general $\overline{T}$ period model
- Point identification of b, p, and costs $F(t)$

Assumption

Nonstandards-track projects have the same cost function $F(t)$ as standards-track projects, but are always published (i.e., $p = b = 1$).

- Second identification strategy, exploiting nonstandards-track RFCs as unique institutional feature of the IETF
- Reasonable assumption: the two tracks are very similar in terms of content and publication process, but differ in commercial implication
- Under this Assumption 2, costs are identified by payoffs and distribution of stopping times for nonstandards-track RFCs, while b and p identified by the payoffs and stopping times for the standards-track RFCs

From Two-Periods Model to ML Estimation

- Time horizon: $\overline{T} = 25$ (robust to higher values of $\overline{T}$)

- Quadratic incremental costs:

$$F(t) = C_0 + C_1 t + C_2 t^2$$

- $\varepsilon_t \sim \text{Logistic}(0, 1)$ with CDF

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

Estimation Procedure

1. Values $\hat{\pi}(t)$ are predicted citations of projects in U.S. patents
2. If Assumption 2: Estimate costs for sample of nonstandards-track projects, based on $\hat{\pi}(t)$ for nonstandards-track from Step 1.
3. Iteratively search for values of $(b, p, \boldsymbol{F})$ that maximize the log-likelihood, retaining estimates of $\boldsymbol{F}$ from Step 2 if using Assumption 2.

Results

Baseline: Ex-Ante Identical Projects

	Full Sample	WG Sample
Rate of Learning (b)	0.17	0.34
	[0.16,0.20]	[0.32,0.37]
Quality Prior (p)	0.59	0.73
	[0.51,0.73]	[0.64,0.83]
Costs $F(1)$	2.35	3.74
	[1.95,3.25]	[2.90,5.15]
Costs $F(10)$	1.25	1.14
	[1.12,1.81]	[1.01,1.72]
Costs $F(20)$	0.51	0.66
	[0.48,0.65]	[0.59,0.86]
Projects: Standards track	14,444	3,168

Empirical and Simulated Hazard Rates

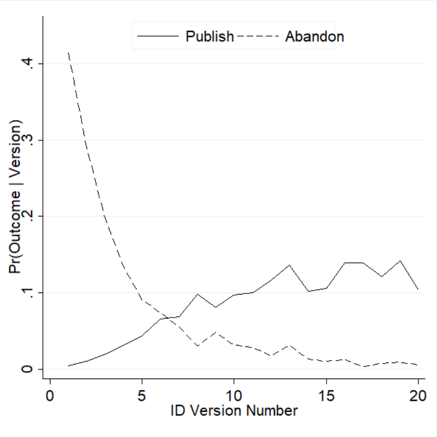

IETF Data

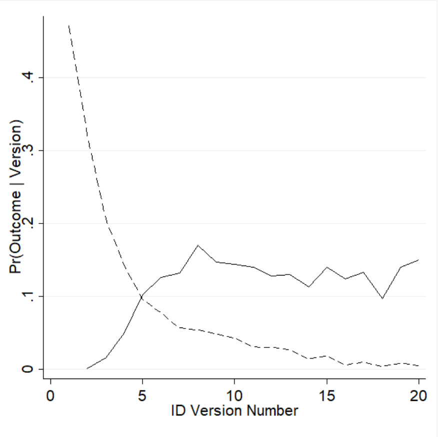

Simulated Data

Robustness to Alternative Assumptions

1. Area Specific Citations:
    - Allow for citations (Step 1) to vary with observed technology areas
2. Calendar Time:
    - Duration of projects measured in quarters rather than versions
3. Phase-Specific Costs:
    - Relax Assumption 1: $F^{\text{post}}(t) = F^{\text{pre}}(t) + \kappa$
4. Three-Step Estimation:
    - Including Assumption 2
5. Others: Include censored projects; RFC citations; $\overline{T} = 50$; only first project for each WG

Robustness Results

	Baseline	Area Specific	Calendar Time	Phase Specific (κ)	3-Step Estimator
Rate of Learning b	0.17	0.10	0.20	0.26	0.30
	[0.16,0.20]	(0.003)	(0.003)	(0.006)	[0.27,0.34]
Quality Prior p	0.59	0.59	0.49	0.42	0.45
	[0.51,0.73]	(0.008)	(0.003)	(0.007)	[0.28,0.59]

Baseline Estimates: Summary

- Roughly 60% of projects submitted to IETF can potentially generate value
- One third of these is published (successful)
- Arrival rate of breakthrough/consensus is 17% per revision
- Our parsimonious model is robust to changes in estimation assumptions and sampling
- Inform theory literature by estimating parameters driving Bayesian learning process in models of experimentation

Counterfactuals

Counterfactual Simulations

- R&D Cost Subsidies and Publication Prizes
  - Which of these generates more value?
- Over-confidence and Pessimism
  - What are the costs of misaligned priors?
  - Is over-confidence good for R&D?

Prizes and Subsidies

- **Prize:** award $\delta_p \geq 0$ “citations” to published RFCs
  - akin to a publication prize, or higher patent values
- **Subsidy:** reduce “citation cost” by $\delta_s \geq 0$ for *all* revisions
  - akin to an R&D tax credit: not conditional on success of the investment
- Keeping the overall budget (in terms of cost of the regime) constant:
  - How do these policies affect authors’ incentives?
  - Which policy generates larger values (to authors/organization)?

Incentive Effects

- Differential effects on exploitation-exploration trade-off
- **Prize:** effect weakens in t
  - After consensus there is no effect, hence no effect on projects that are successful in the baseline without a prize.

$$\text{incremental increase of } \hat{\pi} = \hat{\pi}(t+1) + \delta_p - [\hat{\pi}(t) + \delta_p]$$

  - Before consensus, the relative effect is decreasing with t because authors update beliefs and expect to receive prize with lower probability
- **Subsidy:** effect increases in t
  - Paid per revision, both before and after consensus
  - Constant in absolute terms but increasing in relative terms as incremental opportunity costs are decreasing

Citations (Gross Value)

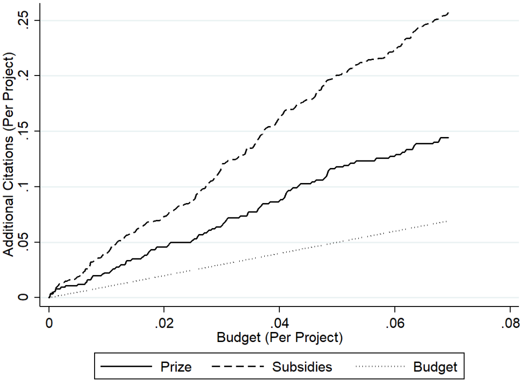

- Subsidies with stronger effect on citations than prize
- Subsidies with effect both on intensive and extensive margin
- Both policies with elasticity larger than one

Net Project Value

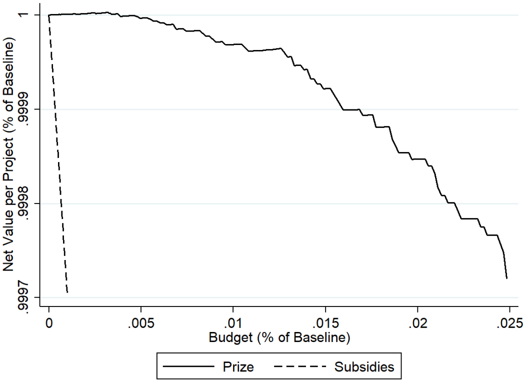

- Total value of projects, accounting for authors' costs of submitting (more) revisions
- Only a (modest) prize can increase value

Optimism and Pessimism

- What are the costs of misaligned priors about project quality?
- Suppose p is the true (estimated) prior quality, and authors have subjective priors p'.
  - True prior determines the selection of project types (good and bad)
  - Subjective priors determine the authors' decisions to continue
- We refer to $p' > p$ as "over-confidence" and and $p' < p$ as "pessimism"
- Take (p, b) as given, and consider an author-team with prior p'; simulate standards-track projects for $p' \in (0, 1)$

Project Value

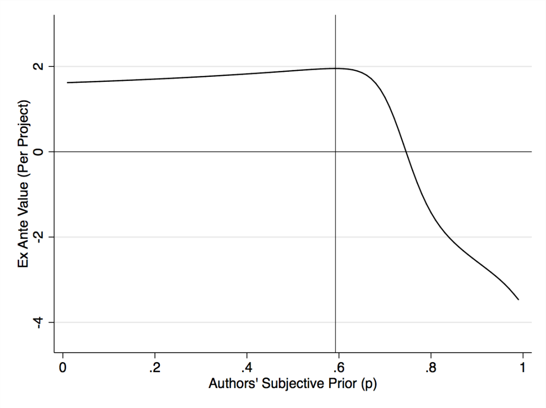

- Both over-confidence and pessimism are distortive
- Costs of over-confidence in terms of total value are higher
- Over-confident researchers over-invest in unpromising lines of research (“dry wells")

Project Duration

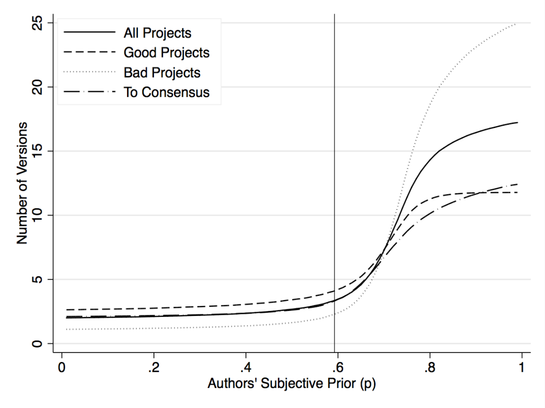

- More over-confident author-teams continue longer
- For higher values of p', this effect is driven by bad projects

Results: Project Heterogeneity

Exploit Observed Project Heterogeneity

1. Communication: positive effect on rate of learning
2. Experience: positive effect on rate of learning
3. Commerciality: suit-to-beard ratio with negative effect on rate of learning
4. Team Size: positive effect on rate of learning

Summary of Results

- Parsimonious model of dynamic learning in R&D decision-making at the project level
- Data from the IETF (with institutional features for identification)
- Learning (b) "about use of time:"
    - Higher rate of learning: bad projects are abandoned faster, yielding higher-value process
    - More communication, experience, commerciality, and team size induce a higher rate of learning
- Prizes yield better outcome than subsidies (~R&D tax credit) when accounting for private costs of overdeveloping bad ideas
- Over-confidence is more costly than pessimism

What's Next?

- What is in a patent?
  - Combine data on patent declarations to the IETF to study how patent disclosure changes project content
- Network structure
  - Consider researchers in a network: Is there a relationship between nature of network ties, author affiliation/centrality/experience and outcomes?
- Attention
  - Use information on emails to study causal relationship between communication and outcomes

Thank you!

Please send comments or suggestions to

ganglmair@utdallas.edu
tsimcoe@bu.edu
or tarantino@uni-mannheim.de

Major Contributors

Count of Project Submissions

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

Dynamic Host Configuration Protocol

RFC 2131

Abstract

The Dynamic Host Configuration Protocol (DHCP) provides a framework for passing configuration information to hosts on a TCPIP network. DHCP is based on the Bootstrap Protocol (BOOTP) [7], adding the capability of automatic allocation of reusable network addresses and additional configuration options [19]. DHCP captures the behavior of BOOTP relay agents [7, 21], and DHCP participants can interoperate with BOOTP participants [9].

Back

Known TCP Implementation Problems

RFC 2525

Introduction

This memo catalogs a number of known TCP implementation problems. The goal in doing so is to improve conditions in the existing Internet by enhancing the quality of current TCP/IP implementations. It is hoped that both performance and correctness issues can be resolved by making implementors aware of the problems and their solutions. In the long term, it is hoped that this will provide a reduction in unnecessary traffic on the network, the rate of connection failures due to protocol errors, and load on network servers due to time spent processing both unsuccessful connections and retransmitted data. This will help to ensure the stability of the global Internet.

Back