

# Lecture

10/26

Mid term Thur  
- opps!

Slides due Nov 4  
- soon!

- then 2 presentations per class
- (I have to do my slides too)
  - what day did we schedule stuff

## Midterm

- anything that happened in class
- go through handout / PPTs
- required reading
- short answer + definition
- what are types of differential privacy
  - + give examples
- multiple choices
- conceptual - inc examples from cases
- 75 min
- no past exam posted
- closed book
- not from blog
  - but good angle to read

(2)

- quickly
- chess is a limited space  $\rightarrow$  math
- Words much harder
  - only grounded in individual understanding
  - b/c humans have common experience
- some things easier for computer
  - calculation  $\checkmark$
  - relational db w/ strict schema  $\checkmark$
  - but when compare Dave  $\rightarrow$  David?
    - need special logic
    - not 100% confident
  - hard in Jeopardy when info in clues is 'implicit'

Broad Domain  
Complex Lang  
High Precision  
Accurate Confidence  
High Speed

Confidence important - since lose \$  
- and give competitors time

Speed important



②

## Presentation

- part of grade → presentation skills
- good feedback for written project

More focused  
More fact-base } Best

- not impressions + opinions

Put forward hypothesis + support

---

## IBM's Watson

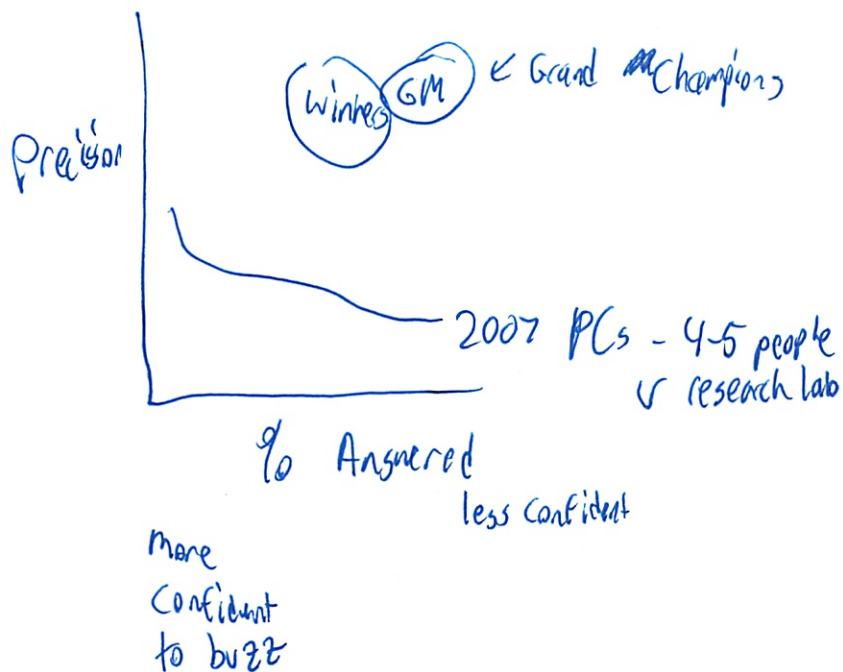
- natural language

## Interview David Ferrucci

- drives scientific advances
- computers impact society + science
- BI, health, e-discovery
- natural language
- imagination of public like w/ Deep Blue
- broad domain of knowledge
  - not a specific database
  - determine what is being asked?
  - find answer
  - how sure am I?

③

Can't build specialized, formal database  
- Very long tail



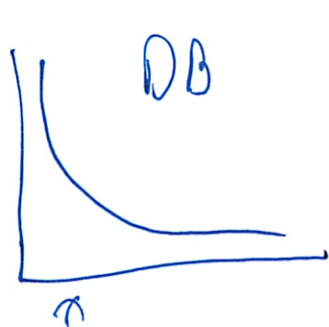
Ken focused on buzzing strategy

- buzzed fast - ans sometimes 70% of the board  
And fast


hand crafted db too narrow

- had to map the data each time
- bottleneck of mapping knowledge into schema
- do some of that
  - very confident on 10% of data
- even medicine + law domain is very large

(4)



narrow +  
precise



lots of coverage  
but not much confidence/  
precision

So know one component

Run many algorithms in parallel

Try interpreting many ways

- massive parallelism is key

do Syntactic frames

- nouns
- verbs
- objects

then probability of each


- w/ different methods

does clue say what type - ie 'liquid'?

- well see Cytoplasm is a fluid
- but did not build 'liquid = fluid'
- Wordnet did build taxonomy
  - returns false

5) But can see in leared text used synonymously

- more trustworthy in Jopardy
- in other domains would not be trusted as much

Cushioning?   
surrounding - depends on context

then identify gaps

get humans to fill gaps

---

can't just brush match keywords

- temporal → 400th anniversary → go back to 1498
- statistical paraphrasing <sup>On May 1898</sup>  
↳ arrived = landed
- geospatial  
↳ Kappad Beach in India

Assess and weigh many algorithms' possible solutions

- in competition with each other to provide effect
- more complex algorithms actually provide less of the ans now vs earlier
- better more general algorithms




6

- if behind AI in game it takes more risks  
need a pun scorer

learns within categories

- changes in confidence in what form ans is in



Competitive w/ human players

- some still beat

Increases getting smaller

- very difficult qv

- humans need broader knowledge

- 1 or 2 years 70% - 80%

- after that a lot more work

- totally new techniques

- diminishing return on Jeopardy

But now turning to applications

Would take 2 hrs on standard desktop

3000 cores in system

Since task is parallel

⑦

More cores  $\rightarrow$  do more, but not faster really  
w/o lots of optimization

Confidence is important  
threshold of betting decision

will not answer even if 90% sure if no reason to take a risk

ranked 1st in half of games w/ champions

---

Useful in medical diagnostic

Q Where use

- health care
- primary care
- don't want to look too much at specific models
- broad base of knowledge
- esp where time constrained

Q Tech support + Ent. knowledge

working w/

can't do popularity

Q Web search

Wolfram Alpha does lots of db

- very different approach

biz model about advertising

- if precise, would you not click around

⑧

Q: How good at controversial ans?

- how old is earth?
- broadening context which you judge evidence
  - which sources are reliable
  - opinion, fiction
  - research area?

Q: What does it read?

- train
- judges source for that domain
- many sources are mixed

Q: How debug when it gets wrong  
drill down into algorithm

all shown in a Java window

see confidence on each sub algorithm

built tool for this

Q: boundaries? How deep - depends on source?

Can it go beyond facts? Opinions?

- Understands what is in context
- how to make it smarter intended knowledge
- what reason that ans not make sense?
- can I encode enough background high level lang
- where can acquire

⑨ Q: In health care not 1 line answers: What hurdles

Yes is a challenge

Better than you think

Summarization + Organization

Coming up w/ new procedures - not practical

- lots of domain expertise

- not trying to solve

Points you at evidence to read

Q: when commercial system?

3-5 years

decision support

~~Q: when~~

- 10-15 years can ask follow up, etc

- need to invest in this

- need to be excited about it

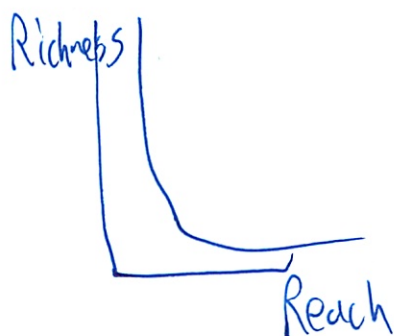



The info I remember fairly well  
 it's the math that will be hard  
 and no pre test!!!

Copy down the key difficulties from slide  
 production cost of info good = 0

$$C(x) = F + m(x)$$

$\uparrow$  total cost       $\uparrow$  fixed       $\uparrow$  marginal




Moore's Law Exponential growth in PC hardware  
 doubling 18 mts

Metacalf's Law quadratic network effect


Differential Pricing

- charge on MC's
- charge of fixed price's
- traditional monopoly  $MR = MC$  leaves half the money on the table

(2)



if can price discriminate



~~Monopoly~~

3 forms

- personalized / 1st degree
  - cor salesman
- Versioning / 2nd degree
  - users choose from product line
  - ie qtd. discounts
- Group Pricing / 3rd degree
  - group membership / identity

Only illegal if lessens competition  
increases economic welfare (did today in 14.01)  
- but not equity

Have we got personalized pricing online?  
- hard to do

- People get angry at (Amazon)

③

## Versioning


- reason train is uncovered
- not cost, but to make it undesirable

Self selection + screening


Online vs offline

## Revenue based pricing

- pay as you go



Upfront fee



revenue based




## Guidelines

- watch out for arbitrage
- aim for complements  
online vs offline
- Start w/ best + value subtract
- When in doubt offer 3 versions

## Group

- student, sr citizens

# ④ Bundling



$$.5 \cdot .5 + .5 \cdot .5$$

$$.5$$

→

$$1 \cdot 1$$

→

$$1$$

\* this is for unrelated/uncorrelated goods

every consumer gets value mix between

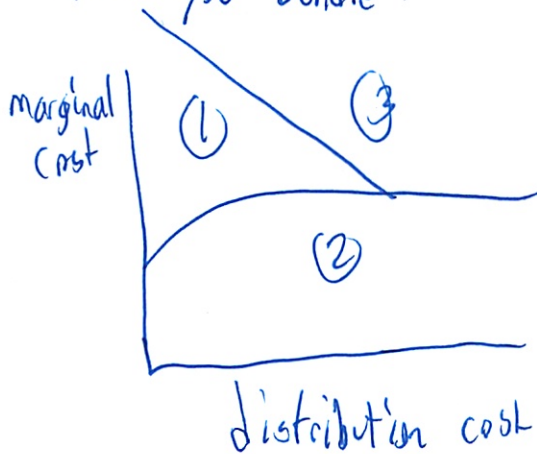
- like some goods

- don't like others

- demand concentrates near center



So should you bundle?



1. unbundle

2. bundle

3. all sales unprofitable

if consumers don't value a lot of the goods the MC becomes magnified



⑤

Site license - aggregate across users

Subscription - aggregate over time

Can aggregate/disaggregate on several dimensions simultaneously

Pitfalls

confusion + resentment

arbitrage - easy to convert

implementation cost

low?

having goods to bundle

need low MC

if valuations vary widely


could offer menu of bundles (differential pricing)

aggregation makes competition tougher

- winner take all

~~works~~ works if sum goods lower quality than others

perfectly negative correlated



bundling doubles revenue

perfectly correlated

bundling no effect


perfectly <sup>vn</sup> correlated - see top pg 4

## ⑥ Search, Competition + Pricing

- is internet great equalizer?
- kinda, but still differences

### Botrand Competition

- oligopoly people buy from the lowest price of the two sellers



eliminating  
search cost eliminates ignorance premium

but wide range of prices online

why - website usability + trust

6-16% lower price online

better price discrimination captures more area under curve

buyers pick lowest price 47%

when using intermediary

service quality, branding, shipping  
not equalizer

## ⑦ Two Sided Networks

give away one side - free adobe reader

gain in profits from install base larger than drop of making  
↓ free

subsidize the one that creates more cross surplus

## Google

(not PPT, so may be missing details)

$\frac{p}{A}$	4	Quality score
B	3	- relevance
C	2	- predicted CTR
D	1	- speed landing pg

Bid · quality score = effective bid

Revenue = click · <sup>effective</sup> bid

price actually pay

$$p = \frac{B_2 Q_2}{Q_1} = \frac{9}{6} \quad \text{€ pays } \$1.50$$

next

~~next~~ pays  $\frac{1.8}{3} = 2 \cdot \frac{2}{3}$  actually pays more

---



## Long Tail

amazon actually meets a lot more consumer surplus  
by providing access to rare books

8

Not a lot more selection, rarer products

estimate power law



$$\ln(Q) = \ln(a) + b \ln(\text{Rank}) + \epsilon$$

$\epsilon$  error

$$a = 37,274$$

$$b = ,8571$$

36% of sales in books not sold in Borders  
much greater surplus in ~~selection~~ selection than if price ↓  
should prices rise or fall on long tail?

depends if need or want

↑ (encourage exploration)



## Talked w/ Erik about Paper

10/28

Wants grounded in reality

"Fine to talk about future"

- total opposite what TIA had us  
and economic theory of class

- this I may be able to add to

- or not me

- but does want discussions of tech to some extent

- not all abstract

- so I think mine is mostly good

- consumer reaction evidence

- pear down

- paper should not focus on tech of course

- will need to see whole thing to adjust

## 15.567 Reading Guide

### October 26: DeepQA and Knowledge Work

The reason we are more productive than any other society in history is not because we work harder, or even because we use more resources. It's because we have more knowledge. Advances in IT, and in particular Artificial Intelligence and Knowledge Management, are transforming our ability create, capture and share knowledge. In this session, we will learn about one such system, IBM's Watson, as well as some of the limitation to knowledge management. We will also discuss a taxonomy for knowledge types, inspired by F.A Hayek's seminar article.

#### Readings

Clive Thompson "What is IBM's Watson?" *New York Times Magazine*, June 16, 2010.

Hayek, F.A. "The Use of Knowledge in Society." *American Economic Review* 35:4 (September 1945): 519-530.

#### Study Questions

Please think about the following issues and questions as you do the readings.

1. What is the distinction between specific vs. general knowledge? Which, if any, does IBM's Watson help with?
2. What other technologies are relevant for knowledge management and knowledge sharing?
3. Which jobs will be affected in next 10 years? What are managerial and organizational implications?
4. What are competitive and strategic implications?
5. In what year will we be able to hand over most important decisions to our robot overlords?

## 15.567 Reading Guide

### November 2: Wikipedia

The Wikipedia is a fascinating example of a collaborative project that has emerged on the Web. How does it work, both the technology and the economics? What advantages and disadvantages does this have? Is it sustainable? What's next? We'll explore these issues in this class.

#### Readings

##### **Required:**

- Wikipedia case (2009), Northwestern University (Greenstein and Devereux)

##### **Optional:**

- Eric S. Raymond (2001). "The Cathedral and the Bazaar" in *The Cathedral & the Bazaar*. O'Reilly.

##### **URLs:**

<http://www.wikipedia.org>.

#### **Study Questions**

Please think about the following questions as you do the readings.

1. Before class, please visit Wikipedia and read a few articles on topics where you are knowledgeable or even an "expert". Make a positive contribution to one article, either by correcting a mistake or providing additional useful detail. For this assignment, your change can be very minor, but do make an actual (and constructive) change of some sort to the Wikipedia. Please come prepared to briefly describe the contribution you made to Wikipedia.
2. What is a wiki and how do they work? What are the similarities and differences between Wikipedia and other websites?
3. What kinds of division of labor have emerged at Wikipedia? How do revision wars get started and settled?
4. What other aspects of Wikipedia etiquette are important?
5. Is this a good way to arrive at truth? Is it a good way to create a useful website?
6. How can businesses learn from Wikipedia? Can Wikis be run on a profit-based model? If so, how? If not, why not?



# Wikipedia:Notability (academics)

From Wikipedia, the free encyclopedia  
(Redirected from Wikipedia:PROFESSOR)

*WP:ACADEMIC* redirects here. For a list of academic studies on Wikipedia, see *Wikipedia:Wikipedia in academic studies*.  
The Deletion Sorting Project has a list of deletion of articles related to Academics and educators  
*Wikipedia:WikiProject Deletion sorting/Academics and educators*



**This page documents an English Wikipedia notability guideline.** It is a generally accepted standard that editors should attempt to follow, though it is best treated with common sense, and occasional exceptions may apply. Any substantive edit to this page should reflect consensus. When in doubt, discuss first on the talk page.

## This page in a nutshell:



- Subjects of biographies are required to be notable; that is significant, interesting, or unusual enough to be worthy of notice, as evidenced by being the subject of significant coverage in independent reliable secondary sources.
- Many scientists, researchers, philosophers and other scholars (collectively referred to as "academics" for convenience) are notably influential in the world of ideas without their biographies being **the subject** of secondary sources.
- This notability guideline specifies criteria for judging the notability of an academic through reliable sources for the impact of their work.

This guideline, sometimes referred to as the **professor test**, is meant to reflect consensus about the notability of academics as measured by their academic achievements. For the purposes of this guideline an academic is someone engaged in scholarly research or higher education and academic notability refers to being *known for such engagement*.<sup>[1]</sup>

This guideline is *independent* from the other subject specific notability guidelines, such as WP:BIO, WP:MUSIC, etc.: it is possible for an academic not to be notable under the provisions of this guideline but to be notable in some other way under one of the other subject specific notability guidelines. Conversely, if an academic is notable under this guideline, his or her possible failure to meet other subject specific notability guidelines is irrelevant.

## Contents

- 1 Criteria
- 2 Notes and examples



- 3 Citation metrics
- 4 Caveats
- 5 Footnotes

## Criteria

If an academic/professor meets *any one* of the following conditions, as substantiated through reliable sources, they are notable. If an academic/professor meets none of these conditions, they *may* still be notable, if they meet the conditions of WP:Notability or other notability criteria, and the merits of an article on the academic/professor will depend largely on the extent to which it is verifiable. See the **Notes and Examples** section below *before applying this guideline*.

1. The person's research has made significant impact in their scholarly discipline, broadly construed, as demonstrated by independent reliable sources.
2. The person has received a highly prestigious academic award or honor at a national or international level.
3. The person is or has been an elected member of a highly selective and prestigious scholarly society or association (e.g. a National Academy of Sciences or the Royal Society) or a Fellow of a major scholarly society for which that is a highly selective honor (e.g. the IEEE)
4. The person's academic work has made a significant impact in the area of higher education, affecting a substantial number of academic institutions.
5. The person holds or has held a named chair appointment or "Distinguished Professor" appointment at a major institution of higher education and research.
6. The person has held a major highest-level elected or appointed academic post at a major academic institution or major academic society.
7. The person has made substantial impact outside academia in their academic capacity.
8. The person is or has been an editor-in-chief of a major well-established journal in their subject area.
9. The person is in a field of literature (e.g writer or poet) or the fine arts (e.g. musician, composer, artist), and meets the standards for notability in that art, such as WP:CREATIVE or WP:MUSIC.

For people who have made substantial impact outside academia but in their academic capacity, the appropriate criteria for that sort of notability apply as an alternative—as for a person notable for popular writing in her subject. If notable only in another capacity entirely, see the general criteria for that field.

An alternative standard, "the academic is more notable than the average college instructor/professor" is often cited. This criterion has the advantage of being concise though it is not universally accepted. Determining the notability of an average professor is difficult in itself and usually relies on one of the nine more detailed criteria above. When used, this criterion is generally applied to indicate that a tenured full or associate professor in a high ranking institution in the US, or equivalent rank elsewhere, is above the average.

It is possible for an academic to be notable according to this standard, and yet not be an appropriate topic for coverage in Wikipedia because of a lack of reliable, independent sources on the subject. Every topic on Wikipedia must be one for which sources exist; see Wikipedia:Verifiability. For the routine uncontroversial details of a career, official institutional and professional sources are accepted.

## Notes and examples

### Notability



Examples and practical tips for applications of this guideline follow.

1. The most typical way of satisfying Criterion 1 is to show that the academic has been an author of *highly cited* academic work:<sup>[2]</sup> either of several extremely highly cited scholarly publications or of a substantial number of scholarly publications with significant citation rates. Reviews of the person's work, published in selective academic publications,<sup>[3]</sup> can be considered together with ordinary citations here. Differences in typical citation and publication rates and in publication conventions between different academic disciplines should be taken into account.<sup>[4]</sup>
2. Criterion 1 can also be satisfied if the person has pioneered or developed a significant new concept, technique or idea, made a significant discovery or solved a major problem in their academic discipline. In this case it is necessary to explicitly demonstrate, by a substantial number of references to academic publications of researchers other than the person in question, that this contribution is indeed widely considered to be significant and is widely attributed to the person in question.
3. There are other considerations that may be used as contributing factors (usually not sufficient individually) towards satisfying Criterion 1, e.g.: significant academic awards and honors (see below); service on editorial boards of scholarly publications; publications in especially prestigious and selective academic journals; publication of collected works; publication of an anniversary or memorial journal volume or a Festschrift dedicated to a particular person; special conferences dedicated to honor academic achievements of a particular person; naming of academic awards or lecture series after a particular person; and others.
4. For the purposes of partially satisfying Criterion 1, significant academic awards and honors may include, for example: major academic awards (they would also automatically satisfy Criterion 2), highly selective fellowships (other than postdoctoral fellowships); invited lectures at meetings of national or international scholarly societies, where giving such an invited lecture is considered considerably more prestigious than giving an invited lecture at typical national and international conferences in that discipline; named lectures or named lecture series; awards by notable academic and scholarly societies; honorary degrees; and others. Ordinary colloquia and seminar talks and invited lectures at scholarly conferences, standard research grants, named post-doctoral fellowships, visiting appointments, or internal university awards are insufficient for this purpose.
5. For the purposes of satisfying Criterion 1, the academic discipline of the person in question needs to be *sufficiently broadly construed*. Major disciplines, such as physics, mathematics, history, political science, or their significant subdisciplines (e.g. particle physics, algebraic geometry, medieval history, "fluid mechanics", "*Drosophila* genetics" are valid examples). Overly narrow and highly specialized categories should be avoided. Arguing that someone is an expert in an extremely narrow area of study is, in and of itself, not necessarily sufficient to satisfy Criterion 1, except for the actual leaders in those subjects.
6. Simply having authored a large number of published academic works *is not considered sufficient to satisfy Criterion 1*.
7. Having a small collaboration distance from a famous or notable academic (e.g. having a small Erdos number) is not, in and of itself, indicative of satisfying Criterion 1.
8. For the purposes of Criterion 2, *major* academic awards, such as the Nobel Prize, MacArthur

#### Subject specific guidelines

##### Academics

Books

Events

Films

Music

Numbers

Organizations & companies

People

Sports and athletes

Web content

##### Active proposals

Periodicals

Fiction

Fatal hull loss civil aviation accidents

##### See also

Common deletion outcomes

Why was my page deleted?

Help (noticeboard)



Fellowship, the Fields Medal, the Bancroft Prize, the Pulitzer Prize for History, etc, *always* qualify under Criterion 2. Some lesser significant academic honors and awards that confer a high level of academic prestige also can be used to satisfy Criterion 2. Examples may include certain awards, honors and prizes of notable academic societies, of notable foundations and trusts (e.g. the Guggenheim Fellowship, Linguapax Prize), etc. Significant academic awards and honors can also be used to partially satisfy Criterion 1 (see item 4 above in this section).

9. Victories in academic student competitions at the high school and university level as well as other awards and honors for academic student achievements (at either high school, undergraduate or graduate level) do not qualify under Criterion 2 and do not count towards partially satisfying Criterion 1.
10. Biographical listings in and awards from vanity press publishers, such as the American Biographical Institute, or from publications incorporating a substantial vanity press element in their business model, such as Marquis Who's Who, do not qualify for satisfying Criterion 2 or for partially satisfying Criterion 1.
11. For the purposes of Criterion 3, elected memberships in minor and non-notable societies are insufficient (most newly formed societies fall into that category).
12. Criterion 4 may be satisfied, for example, if the person has authored several books that are widely used as textbooks (or as a basis for a course) at multiple institutions of higher education.
13. Criterion 6 may be satisfied, for example, if the person has held the post of President or Chancellor (or Vice-Chancellor in countries where this is the top academic post) of a significant accredited college or university, director of a highly regarded notable academic independent research institute or center (which is not a part of a university), president of a notable national or international scholarly society, etc. Lesser administrative posts (Provost, Dean, Department Chair, etc) are generally not sufficient to qualify under Criterion 6 alone, although exceptions are possible on a case-by-case basis (e.g. being a Provost of a major university may sometimes qualify). Heads of institutes and centers devoted to promoting pseudo-science and marginal or fringe theories are generally not covered by Criterion 6; their heads may still be notable under other criteria of this guideline or under the general WP:BIO or WP:N guidelines.
14. Criterion 7 may be satisfied, for example, if the person is frequently quoted in conventional media as an academic expert in a particular area. A small number of quotations, especially in local news media, is not unexpected for academics and so falls short of this mark.
15. Criterion 7 may also be satisfied if the person has authored widely popular general audience books on academic subjects provided the author is widely regarded inside academia as a well-established academic expert and provided the books deal with that expert's field of study. Books on pseudo-science and marginal or fringe scientific theories are generally not covered by this criterion; their authors may still be notable under other criteria of this guideline or under the general WP:BIO or WP:N guidelines.
16. Patents, commercial and financial applications are generally not indicative of satisfying Criterion 7.
17. Journals dedicated to promoting pseudo-science and marginal or fringe theories are generally not covered by Criterion 8; their Editors-in-Chief may still be notable under other criteria of this guideline or under the general WP:BIO or WP:N guidelines.

## Citation metrics

The only reasonably accurate way of finding citations to journal articles in most subjects is to use one of the two major citation indexes, Web of Knowledge and Scopus. They are, unfortunately, very expensive: Scopus will be found mostly in university and large college libraries, and Web of Knowledge in major universities. Scopus covers the sciences and the social sciences, but is very incomplete before 1996; Web of Knowledge *may* cover the sciences back to 1900, the social sciences back to 1956, and the humanities back to 1975, but only the largest universities can afford the entire set. (Fortunately, additional citation indexes



with public access are being developed.) These databases are furthermore incomplete especially for the less developed countries. Additionally, they list citations only from journal articles—citations from articles published in books or other publications are not included. For that reason, these databases should not be used for disciplines such as computer science in which conference or other non-journal publication is essential.<sup>[5]</sup> Web of Knowledge provides a free index of highly cited researchers (<http://hcr3.isiknowledge.com/formBrowse.cgi>), which may be of some value. In individual scientific fields, MathSciNet, SciFinder Scholar (Chemical Abstracts), and similar disciplinary indexes are also valuable resources, often specifically listing citation counts, but access to them is also not free and usually requires a university computer account.

- A caution about Google Scholar: Google Scholar works well for fields where all (or nearly all) respected venues have an online presence. Most papers written by a computer scientist will show up, but for less technologically up-to-date fields, it is dicey. For non-scientific subjects, it is especially dicey. Even the journal *Science* puts articles online only back to 1996. Many journals, additionally, do not permit Google Scholar to list their articles. For books, the coverage in Google Scholar is partly through Google Book Search, and is very strongly influenced by publisher's permissions and policies. Thus, the absence of references in Google Scholar should not be used as proof of non-notability. In the other direction, GS includes sources that are not peer-reviewed, such as academic web sites. Thus, the number of citations found there can sometimes be twice the number of actual citations from truly reliable scholarly material. In essence, it is a rough guide only.
- A caution about PubMed: Medline, now usually accessed as part of PubMed, is a well-established broadly based search engine, covering much of biology and all of medicine, published since 1967 and sometimes even earlier. It includes a few journals in medically related clinical subjects, but is not complete in those. Further, not all articles in PubMed are from peer-reviewed journals, as it includes medical news sources of various degrees of quality, including such items in peer-reviewed journals it does cover. It also exhaustively covers letters to the editor and similar material, not all of which is of any significance.
- A caution about "related articles" In PubMed, and most other databases, "Related articles" are **not** articles that necessarily cite the original; they are articles on the same general topic, usually selected by having title words or citations in common. Some *may* cite the original (and some clearly do not, for they will have been published before the articles in question). They are useful for finding additional papers on a subject, which is the purpose for which they were designed. The only way to count citations using such a listing in, for example, PubMed, is the tedious method of looking at every one of the related articles published after the article in question, finding its "cited article" display, and check if it is there. (Some PubMed records do not list cited articles, for a variety of reasons.) Nor will such a listing necessarily include all the citations. - Help for "Related articles" feature ([http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Search&db=books&doptcmdl=GenBookHL&term=%22related+articles%22+AND+helppubmed%5Bbook%5D+AND+404127%5Buid%5D&rid=helppubmed.section.pubmedhelp.Searching\\_PubMed#pubmedhelp.Finding\\_articles\\_rel|PubMed](http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Search&db=books&doptcmdl=GenBookHL&term=%22related+articles%22+AND+helppubmed%5Bbook%5D+AND+404127%5Buid%5D&rid=helppubmed.section.pubmedhelp.Searching_PubMed#pubmedhelp.Finding_articles_rel|PubMed))
- Measures of citability such as H-index, G-index, etc, may be used as a rough guide in evaluating whether Criterion 1 is satisfied, but they should be approached with considerable caution since their validity is not, at present, widely accepted, and since they depend substantially on the source indices used.
- For scholars in humanities the existing citation indices and GoogleScholar often provide inadequate and incomplete information. In these cases one can also look at how widely the person's books are held in various academic libraries (this information is available in Worldcat (<http://www.worldcat.org>) ) when evaluating whether Criterion 1 is satisfied.

## Caveats



Some caveats to this guideline follow.

1. Note that as this is a **guideline** and not a **rule**, exceptions may well exist. Some academics may not meet any of these criteria, but may still be notable for their academic work. It is important to note that it is very difficult to make clear requirements in terms of numbers of publications or their quality: the criteria, in practice, vary greatly by field. Also, this proposal sets the bar fairly low, which is natural: to a degree, academics live in the public arena, trying to influence others with their ideas. It is natural that successful ones should be considered notable.
2. An academic who is not notable by these guidelines could still be notable for non-academic reasons.
3. It is possible for an academic to be notable according to this standard, and yet not be an appropriate topic for an article in Wikipedia because of a lack of reliable, independent sources on the subject. Every topic on Wikipedia must be one for which sources exist; see Wikipedia:Verifiability.

## Footnotes

1. ^ Most academics are or have been faculty members (professors) at colleges or universities. Also, many academics hold or have held academic or research positions in various academic research institutes (such as NIH, CNRS, etc). However, academics, in the sense of above definition, *may also work outside academia* (e.g. in industry, financial sector, government, as a clinical physician, as a practicing lawyer, etc) and their primary job *does not* have to be academic in nature if they are known for their academic achievements; conversely, if they are notable for their primary job, they do not have to be notable academics to warrant an article. School teachers at the secondary education level, sometimes also called professors, are not presumed to be academics and may only be considered academics for the purposes of this guideline if they are engaged in substantial scholarly research and are known for such research. They are rather evaluated by the usual rules for notability in their profession. See professor for more information about academic ranks and their meanings. Note that academic ranks are different in different countries.
2. ^ To count towards satisfying Criterion 1, citations need to occur in peer-reviewed scholarly publications such as journals or academic books.
3. ^ In some disciplines there are review publications that review virtually all refereed publications in that discipline. For example, in mathematics, Mathematical Reviews, also known as MathSciNet, and Zentralblatt MATH fall into that category. The mere fact that an article or a book is reviewed in such a publication does not serve towards satisfying Criterion 1. However, the content of the review and any evaluative comments made there may be used for that purpose.
4. ^ Generally, more experimental and applied subjects tend to have higher publication and citation rates than more theoretical ones. Publication and citation rates in humanities are generally slower than in sciences. Also, in sciences most of new original research is published in journals and conference proceedings whereas in humanities book publications tend to play a larger role. The meaning of "substantial number of publications" and "high citation rates" is to be interpreted in line with the interpretations used by major research institutions in the awarding of tenure.
5. ^ A report from the association of European computer science departments lists ten bullet points for evaluation of computer science research, two of which emphasize the importance of non-journal publication and one of which specifically cautions against the use of Web of Science: Meyer, Bertrand; Choppy, Christine; Staunstrup, Jørgen; van Leeuwen, Jan (2009), "Research Evaluation for Computer Science" (<http://www.informatics-europe.org/docs/research-eval.php>) , *Communications of the ACM* **52** (4): 31–34, doi:10.1145/1498765.1498780 (<http://dx.doi.org/10.1145/2F1498765.1498780>) , <http://www.informatics-europe.org/docs/research-eval.php>. Instead, it recommends Google scholar or CiteSeer for this field.

Retrieved from "[http://en.wikipedia.org/wiki/Wikipedia:Notability\\_\(academics\)](http://en.wikipedia.org/wiki/Wikipedia:Notability_(academics))"

Categories: Wikipedia notability guidelines

- 
- This page was last modified on 13 October 2010 at 01:05.
  - Text is available under the Creative Commons Attribution-ShareAlike License; additional terms may apply. See Terms of Use for details.
- Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.



## Lecture: Wikipedia

11/12

Mean 79

Median 82.5

St. Dev = 13

most common 81-90

question 8 was hardest

50% got

↳ combining bundling + versioning  
"mixed bundling"

---

Feedback: People like guests

Cold calling

Least happy w/ gip project

Do use reading guide

Watson lecture least possible

---

## Watson Debrief

Q: implications for business;

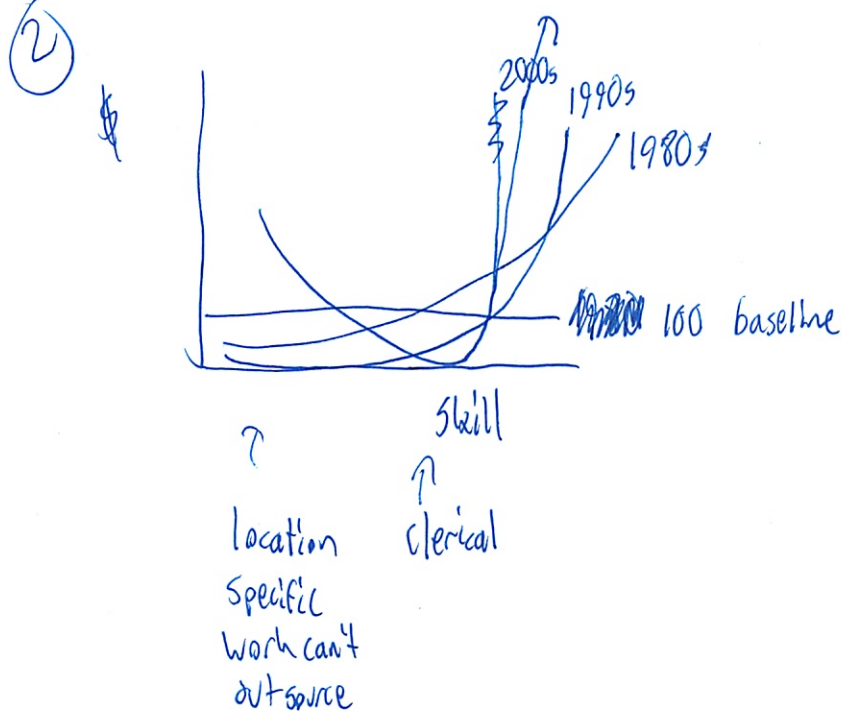
most valuable people were those who knew the people for the process

intranet search is terrible

- ability to think out of box will be more valuable

↳ General knowledge

- needs to add new knowledge



## Wikipedia

meta knowledge - how does knowledge work for society

We are wealthier than past societies since have more knowledge

~~know~~ non rivalrous - can take via making a copy

- harder to monetize
- hard to sell - if you describe it to buyer
  - they have it!
- how to prove you have knowledge w/o giving it
  - 0 sum knowledge proofs

Corp wikis

- some groups use it more than others
- some are else mentioned - Q + A



③

Someone ~~App~~ said Typos are put in to use user engagement

(Something I really know, so write less)

33,000 are the active users

~70% of edits

companies donate into  
- well GDFL

---

Norms No original content ← opposit of academic journals

Cite

Neutral POV

GNU FDL

Only free copyright

Focus on consensus

Attention based on interest

---

(class ~~not~~ missing point)

Should you cite

---

human synthesized vs Google

Consensus  
most of the time

Find ~~the~~ each person's opinion

## Ways to motivate people

- contribute to topics that people are interested in

- need to draw on that knowledge

- could add an extra digit to salary online

top chess team:  $\underbrace{\text{good computers} + \text{good humans}}_{\text{with good process}}$

## 15.567 Reading Guide

### November 4: Enterprise 2.0

The new Web 2.0 technologies are spawning new ways to work inside organizations: Enterprise 2.0. For this class, we will look closely at one cluster of organizations in the midst of this transformation, namely the U.S. Intelligence Community.

#### Readings

##### **Required:**

McAfee, A. The U.S. Intelligence Community (A), MIT Center for Digital Business Case

##### **Optional:**

McAfee, A. "Enterprise 2.0," *Sloan Management Review*, 2006.

Brynjolfsson and McAfee "Beyond Enterprise 2.0", *Sloan Management Review*, Spring, 2007.

##### **URLs:**

<http://andrewmcafee.org/blog/>

<http://www.blogger.com/start>

<http://blogsearch.google.com/>

<http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>

#### **Study Questions**

Please think about the following questions as you do the readings.

- 1) What types of knowledge are there to manage in the U.S. intelligence community? How about at a strategy consulting firm or investment bank? At other types of firms?
- 2) How would you address the knowledge and information-sharing challenges facing the Intelligence Community at the time of the (A) case? What role, if any, would information technology play in your solution?
- 3) Knowledge management is predicated on people sharing information. Why don't people share information? What can you do to encourage knowledge sharing? How have other companies solved this problem?
- 4) What is the role of technology in knowledge management? What technologies are useful for what types of knowledge? Can technology help at all with tacit knowledge? If tacit knowledge is essentially that which cannot be written down, how can any technology help make it more visible?
- 5) Does the US IC have the same knowledge management challenges as, say, HP or Otis Elevator? Why or why not?

# Lecture Enterprise IT

1/6

(did not have reading material)

Enterprise 2.0

- inside orgs

dire failure of gov it

Usually hard to read

~~the~~ intel community does not share info

got very good at collecting

- too much info

agencies protective

Need to know - was watch word

- ↳ turf protection

Each agency has a different domain

More coordination needed than before

(IA does zone defence

FBI goes man to man

no one asked

- all siloed

does not centralized

no connection

must look in the hub



(2)

hubs exist - but not useful

hard to retrieve info

can't connect different pieces of info

No incentives to share info

or poor

pull model

how would push work?

- Some sort of keep me updated

~~is the~~

duplicative efforts

Why design something like this?

Each Agency designed for a different threat separately

- not to work together

What are limitations?

Operating in the US legal system unfair

Afterwards obvious

- but can do it before?

How much did they know?

④③

tlada leads - were on track

Was there unusually high warnings  
- or fade into noise

(Counter) intel threat

- always deeply believed risk of sharing too much info

---

Happens in all firms as they grow

My DB experience

Can't ~~integrate~~ expect too 'integrated + standardized

- too much part of growth

- But too much standardization?

Gravity of being there very hard to overcome

---

Fixes?

council of people w/ authority over other orgs

or decides who knows to know what

Q: how would they stay on top of that?

Split DB - not 1 or hackable

tell what people are working on

~~the~~ db of what each person is working on

---

More inter-departmental meetings

9  
rotational programs of leaders in training

look where info flows

people not aware of what is out there

Info Architect: this is very hard

Get a bit of relation in search

Can you ~~the~~ even do this?

Worked with USSR

Will it work w/ new threats

Dynamic

Who you hire

Opinion of fix outline

- box at top of org chart
- need to multiplex all this data
- start from open model + share as much as possible
- assume people are trustworthy

Read Case port B

⑤

Like people making connections

Need some minimal incentives to contribute

Very efficient ways of capturing value

Clunky UI tools

How do you know people are adding 'right' info?

- quality control

Will I have the time to add stuff to the wiki?

No mention of any org change

- just release the tech

Someone who used it 2 years ago; helps - but

was missing some things

Cultural + policy

Need to ~~add~~ automate some things

---

Can you find stuff on web?

Same characteristics!

- each site designed separately

- organic, unplanned growth

- no center

- always happens w/ growth



(6)

but why?

Web is for sharing

- IC does not have that

lets assume all used it

- how much would that solve their problem?

- ~~the~~ before 9/11 people jumping up + down

- but if leaders still ignore

- know whats most important

Does web let you know what's trending?

- mean

Danger if no one talking about it at all

Intellapedia good at basic updating

How does searching over small data

- search on this site does not matter

Job is to find unknown

Does web help people synthesize info

Democratizing stove pipe

- community decides ~~there~~ what important

Certain communities are more/less active

Community needs to buy in

(can they do it)

- old school, not quick to adapt

⑦ He thinks if everyone did it would be good  
"integration w/o reorg"

---

Since case

- mixed

- tools are known about it

pockets of enthusiasm

islands of indifference + skepticism

mostly managers

Managers can do a lot to prevent this

slow progress

Community has grown

now 800,000

lots of duplication

Staply on the rise

tech there if they wanted

but too much inertia + indifference

Can it change?

**15.567 Reading Guide**  
**November 9: Open Source and Innovation**

Digital goods like software, information, music, video, e-books, etc. can be reproduced almost costlessly. While this provides obvious economic benefits, it also poses special challenges for providing incentives for innovation. Randy Davis discusses some of these challenges in his article on the "Digital Dilemma", while Eric Raymond describes one solution, the open source approach. Understanding these issues will become increasingly important for entrepreneurs, managers, policymakers and innovators of all types.

**Readings**

**Required:**

Raymond, "The Magic Cauldron"

Davis, Randall "The Digital Dilemma" *Communications of the ACM*, February, 2001.

**Optional:**

Shapiro & Varian, Ch. 4

Brynjolfsson and Saunders, Ch. 6.

Richard Stallman, "Why software should not have owners" in *Free Software, Free Society: The Selected Essays of Richard M. Stallman*.

Summary of the Digital Millennium Copyright Act by U.S. Copyright office

Besen and Raskind. (1991). An Introduction to the Law and Economics of Intellectual Property. *Journal of Economic Perspectives*, Vol. 5 (No. 1), pp.3-27.

Lessig, L., *Free Culture: How Big Media Uses Technology and the Law to Lock Down Creativity*, (2004), Penguin Press.

**URLs:**

<http://www.socialtext.com/blog/author/eugene-lee/>

[http://en.wikipedia.org/wiki/Open\\_source](http://en.wikipedia.org/wiki/Open_source)

<http://www.fsf.org>

<http://www.redhat.com/>

## Study Questions

Please think about the following questions as you do the readings.

- 1) What are the strengths and weaknesses of the open source approach, including the business models described by Eric Raymond in "The Magic Cauldron"? Where does he argue that it works best? Do you agree?
- 2) Give your answer to question one, do you think the Open Source approach, is likely to dominate software development in the next 20 years, to fade away, or to stay where it is? Provide specific examples, preferably from your first-hand experience.
- 3) What does Randy Davis mean by the "Digital Dilemma"? What are the trade-offs he discusses? Do you agree with his assessment? Are there other factors that should be considered?
- 4) Should the current laws regarding intellectual property rights, including patents and copyrights be fundamentally revised for software and other digital goods, along the lines suggested by Randall Davis?
- 5) What do you think the best way is to provide incentives for innovation in digital goods like software and media? If had a magic wand, how would you design the "rules of the game" in our economy to maximize welfare in this area?



# Digital Dilemma Reading

11/5

Music biz by Randall Davis

really like argument that's it's a service

↳ speed, reliability, ease of access

not product

↳ since it's not a product - it's a license

↳ so don't bother w/ all these protections

do they work?

do they even help?

right biz models better against individuals

like his POU

- kinda we'd stated

War written in 2001 - he is way ahead of game!

I like his incentive-law proposal

- but what specifically is it?

# The Magic Cauldron

by Eric S. Raymond

June 1999

This paper analyzes the evolving economic substrate of the open-source phenomenon. We first explode some prevalent myths about the funding of program development and the price structure of software. We present a game-theory analysis of the stability of open-source cooperation. We present nine models for sustainable funding of open-source development; two non-profit, seven for-profit. We continue to develop a qualitative theory of when it is economically rational to be closed. We then examine some novel additional mechanisms the market is now inventing to fund for-profit open-source development, including the reinvention of the patronage system and task markets. We conclude with some tentative predictions of the future.

## Contents

1	Indistinguishable From Magic	2
2	Beyond Geeks Bearing Gifts	2
3	The Manufacturing Delusion	3
4	The "information wants to be free" Myth	5
5	The Inverse Commons	6
6	Reasons for Closing Source	7
7	Use-Value Funding Models	8
7.1	The Apache case: cost-sharing	8
7.2	The Cisco case: risk-spreading	9
8	Why Sale Value is Problematic	9
9	Indirect Sale-Value Models	10
9.1	Loss-Leader/Market Positioner	11
9.2	Widget Frosting	11
9.3	Give Away the Recipe, Open A Restaurant	12
9.4	Accessorizing	13
9.5	Free the Future, Sell the Present	13
9.6	Free the Software, Sell the Brand	13
9.7	Free the Software, Sell the Content	13
10	When To Be Open, When To Be Closed	14
10.1	What Are the Payoffs?	14

Seems old CVMugon

from Latex use

↓ same cutoff



why no serious reading

on biz using open source

## 1 Indistinguishable From Magic

In Welsh myth, the goddess Ceridwen owned a great cauldron which would magically produce nourishing food – when commanded by a spell known only to the goddess. In modern science, Buckminster Fuller gave us the concept of ‘ephemeralization’, technology becoming both more effective and less expensive as the physical resources invested in early designs are replaced by more and more information content. Arthur C. Clarke connected the two by observing that “Any sufficiently advanced technology is indistinguishable from magic”.

To many people, the successes of the open-source community seem like an implausible form of magic. High-quality software materializes “for free”, which is nice while it lasts but hardly seems sustainable in the real world of competition and scarce resources. What’s the catch? Is Ceridwen’s cauldron just a conjuring trick? And if not, how does ephemeralization work in this context – what spell is the goddess speaking?

Everyone builds on each others works

## 2 Beyond Geeks Bearing Gifts

The experience of the open-source culture has certainly confounded many of the assumptions of people who learned about software development outside it. “The Cathedral and the Bazaar” 16 () described the ways in which decentralized cooperative software development effectively overturns Brooks’s Law, leading to unprecedented levels of reliability and quality on individual projects. “Homesteading the Noosphere” 16 () examined the social dynamics within which this ‘bazaar’ style of development is situated, arguing that it is most effectively understood not in conventional exchange-economy terms but as what anthropologists call a ‘gift culture’ in which members compete for status by giving things away. In this paper we shall begin by exploding some common myths about software production economics; then continue the analysis of 16 () and 16 () into the realm of economics, game theory and business models, developing new conceptual tools needed to understand the way that the gift culture of open-source developers can sustain itself in an exchange economy.

In order to pursue this line of analysis without distraction, we’ll need to abandon (or at least agree to temporarily ignore) the ‘gift culture’ level of explanation. 16 () posited that gift culture behavior arises in situations where survival goods are abundant enough to make the exchange game no longer very interesting; but while this appears sufficiently powerful as a *psychological* explanation of behavior, it lacks sufficiency as an explanation of the mixed *economic* context in which most open-source developers actually operate. For most, the exchange game has lost its appeal but not its power to constrain. Their behavior has to make sufficient material-scarcity-economics sense to keep them in a gift-culture-supporting zone of surplus.

Therefore, we now will consider (from entirely *within* the realm of scarcity economics) the modes of cooperation and exchange that sustain open-source development. While doing so we will answer the pragmatic question “How do I make money at this?”, in detail and with examples. First, though, we will show that much of the tension behind that question derives from prevailing folk models of software-production economics that are false to fact.

(A final note before the exposition: the discussion and advocacy of open-source development in this paper should not be construed as a case that closed-source development is intrinsically wrong, nor as a brief against intellectual-property rights in software, nor as an altruistic appeal to ‘share’. While these arguments are still beloved of a vocal minority in the open-source development community, experience since 16 () has made it clear that they are unnecessary. An entirely sufficient case for open-source development rests on its engineering and economic outcomes – better quality, higher reliability, lower costs, and increased choice.)

is it a minority

Stallman

oh never mind on the old curriculum bit

do I criticize Eric for going too biz and too tech?



### 3 The Manufacturing Delusion

We need to begin by noticing that computer programs like all other kinds of tools or capital goods, have two distinct kinds of economic value. They have *use value* and *sale value*.

The *use value* of a program is its economic value as a tool. The *sale value* of a program is its value as a salable commodity. (In professional economist-speak, sale value is value as a final good, and use value is value as an intermediate good.)

When most people try to reason about software-production economics, they tend to assume a 'factory model' that is founded on the following fundamental premises.

1. Most developer time is paid for by sale value.
2. The sale value of software is proportional to its development cost (i.e. the cost of the resources required to functionally replicate it) and to its use value.

In other words, people have a strong tendency to assume that software has the value characteristics of a typical manufactured good. But both of these assumptions are demonstrably false.

First, code written for sale is only the tip of the programming iceberg. In the pre-microcomputer era it used to be a commonplace that 90% of all the code in the world was written in-house at banks and insurance companies. This is probably no longer the case – other industries are much more software-intensive now, and the finance industry's share of the total has accordingly dropped – but we'll see shortly that there is empirical evidence that around 95% of code is still written in-house.

This code includes most of the stuff of MIS, the financial- and database-software customizations every medium and large company needs. It includes technical-specialist code like device drivers (almost nobody makes money selling device drivers, a point we'll return to later on). It includes all kinds of embedded code for our increasingly microchip-driven machines - from machine tools and jet airliners to cars to microwave ovens and toasters.

Most such in-house code is integrated with its environment in ways that make reusing or copying it very difficult. (This is true whether the 'environment' is a business office's set of procedures or the fuel-injection system of a combine harvester.) Thus, as the environment changes, there is a lot of work continually needed to keep the software in step.

This is called 'maintenance', and any software engineer or systems analyst will tell you that it makes up the vast majority (more than 75%) of what programmers get paid to do. Accordingly, most programmer-hours are spent (and most programmer salaries are paid for) writing or maintaining in-house code that has no sale value at all – a fact the reader may readily check by examining the listings of programming jobs in any newspaper with a 'Help Wanted' section.

Scanning the employment section of your local newspaper is an enlightening experiment which I urge the reader to perform for him- or herself. Examine the jobs listings under programming, data processing, and software engineering for positions that involve the development of software. Categorize each such job according to whether the software is being developed for use or for sale.

It will quickly become clear that, even given the most inclusive definition of 'for sale', at least nineteen in twenty of the salaries offered are being funded strictly by use-value (that is, value as an intermediate good). This is our reason for believing that only 5% of the industry is sale-value-driven. Note, however, that the rest of the analysis in this paper is relatively insensitive to this number; if it were 15% or even 20%, the economic consequences would remain essentially the same.



(When I speak at technical conferences, I usually begin my talk by asking two questions: how many in the audience are paid to write software, and for how many do their salaries depend on the sale value of software. I generally get a forest of hands for the first question, few or none for the second, and considerable audience surprise at the proportion.)

Second, the theory that the sale value of software is coupled to its development or replacement costs is even more easily demolished by examining the actual behavior of consumers. There are many goods for which a proportion of this kind actually holds (before depreciation) – food, cars, machine tools. There are even many intangible goods for which sale value couples strongly to development and replacement cost – rights to reproduce music or maps or databases, for example. Such goods may retain or even increase their sale value after their original vendor is gone.

By contrast, when a software product's vendor goes out of business (or if the product is merely discontinued), the maximum price consumers will pay for it rapidly goes to near zero regardless of its theoretical use value or the development cost of a functional equivalent. (To check this assertion, examine the remainder bins at any software store near you.)

The behavior of retailers when a vendor folds is very revealing. It tells us that they know something the vendors don't. What they know is this: the price a consumer will pay is effectively capped by the expected future value of vendor service (where 'service' is here construed broadly to include enhancements, upgrades, and follow-on projects).

In other words, software is largely a service industry operating under the persistent but unfounded delusion that it is a manufacturing industry.

It is worth examining why we normally tend to believe otherwise. It may simply be because the small portion of the software industry that manufactures for sale is also the only part that advertises its product. Also, some of the most visible and heavily advertised products are ephemera like games that have little in the way of continuing service requirements (the exception, rather than the rule) 16 ().

It is also worth noting that the manufacturing delusion encourages price structures that are pathologically out of line with the actual breakdown of development costs. If (as is generally accepted) over 75% of a typical software project's life-cycle costs will be in maintenance and debugging and extensions, then the common price policy of charging a high fixed purchase price and relatively low or zero support fees is bound to lead to results that serve all parties poorly.

Consumers lose because, even though software is a service industry, the incentives in the factory model all cut against a vendor's offering competent service. If the vendor's money comes from selling bits, most effort will go to making bits and shoving them out the door; the help desk, not a profit center, will become a dumping ground for the least effective and get only enough resources to avoid actively alienating a critical number of customers.

The other side of this coin is that most vendors buying this factory model will also fail in the longer run. Funding indefinitely-continuing support expenses from a fixed price is only viable in a market that is expanding fast enough to cover the support and life-cycle costs entailed in yesterday's sales with tomorrow's revenues. Once a market matures and sales slow down, most vendors will have no choice but to cut expenses by orphaning the product.

Whether this is done explicitly (by discontinuing the product) or implicitly (by making support hard to get), it has the effect of driving customers to competitors (because it destroys the product's expected future value, which is contingent on that service). In the short run, one can escape this trap by making bug-fix releases pose as new products with a new price attached, but consumers quickly tire of this. In the long run,

where open does not work: Nokia, Marmos

actually I really like the 2 papers he chose: good on biz + tech (via weird explanation)

almost every open source dev has a full time job - best ones paid to maintain the open source

esp corp

hmm never thought of that

does that happen? - oh open source is reverse

- open source good at this

- free upgrades fast to abandon old versions



therefore, the only way to escape is to have no competitors – that is, to have an effective monopoly on one’s market. In the end, there can be only one.

And, indeed, we have repeatedly seen this support-starvation failure mode kill off even strong second-place competitors in a market niche. (The pattern should be particularly clear to anyone who has ever surveyed the history of proprietary PC operating systems, word processors, accounting programs or business software in general.) The perverse incentives set up by the factory model lead to a winner-take-all market dynamic in which even the winner’s customers end up losing.

If not the factory model, then what? To handle the real cost structure of the software life-cycle efficiently (in both the informal and economics-jargon senses of ‘efficiency’), we require a price structure founded on service contracts, subscriptions, and a continuing exchange of value between vendor and customer. Under the efficiency-seeking conditions of the free market, therefore, we can predict that this is the sort of price structure most of a mature software industry will ultimately follow. *and of course biz like that better*

The foregoing begins to give us some insight into why open-source software increasingly poses not merely a technological but an economic challenge to the prevailing order. The effect of making software ‘free’, it seems, is to force us into that service-fee-dominated world – and to expose what a relatively weak prop the sale value of closed-source bits was all along.

The term ‘free’ is misleading in another way as well. Lowering the cost of a good tends to increase, rather than decrease, total investment in the infrastructure that sustains it. When the price of cars goes down, the demand for auto mechanics goes up – which is why even those 5% of programmers now compensated by sale-value would be unlikely to suffer in an open-source world. The people who lose in the transition won’t be programmers, they will be investors who have bet on closed-source strategies where they’re not appropriate.

*like the articles I used to write*

## 4 The “information wants to be free” Myth

There is another myth, equal and opposite to the factory-model delusion, which often confuses peoples’ thinking about the economics of open-source software. It is that “information wants to be free”. This usually unpacks to a claim that the zero marginal cost of reproducing digital information implies that its clearing price ought to be zero.

The most general form of this myth is readily exploded by considering the value of information that constitutes a claim on a rivalrous good – a treasure map, say, or a Swiss bank account number, or a claim on services such as a computer account password. Even though the claiming information can be duplicated at zero cost, the item being claimed cannot be. Hence, the non-zero marginal cost for the item can be inherited by the claiming information.

We mention this myth mainly to assert that it is unrelated to the economic-utility arguments for open source; as we’ll see later, those would generally hold up well even under the assumption that software actually *does* have the (nonzero) value structure of a manufactured good. We therefore have no need to tackle the question of whether software ‘should’ be free or not.



## 5 The Inverse Commons

Having cast a skeptical eye on one prevailing model, let's see if we can build another – a hard-nosed economic explanation of what makes open-source cooperation sustainable.

This is a question that bears examination on a couple of different levels. On one level, we need to explain the behavior of individuals who contribute to open-source projects; on another, we need to understand the economic forces that sustain cooperation on open-source projects like Linux or Apache.

Again, we must first demolish a widespread folk model that interferes with understanding. Over every attempt to explain cooperative behavior there looms the shadow of Garret Hardin's Tragedy of the Commons.

Hardin famously asks us to imagine a green held in common by a village of peasants, who graze their cattle there. But grazing degrades the commons, tearing up grass and leaving muddy patches, which re-grow their cover only slowly. If there is no agreed-on (and enforced!) policy to allocate grazing rights that prevents overgrazing, all parties' incentives push them to run as many cattle as quickly as possible, trying to extract maximum value before the commons degrades into a sea of mud. *actually some of best projects controlled*

Most people have an intuitive model of cooperative behavior that goes much like this. It's not actually a good diagnosis of the economic problems of open-source, which are free-rider (underprovision) rather than congested-public-good (overuse). Nevertheless, it is the analogy I hear behind most off-the-cuff objections. *Firefox Ubuntu*

The tragedy of the commons predicts only three possible outcomes. One is the sea of mud. Another is for some actor with coercive power to enforce an allocation policy on behalf of the village (the communist solution). The third is for the commons to break up as village members fence off bits they can defend and manage sustainably (the property-rights solution).

When people reflexively apply this model to open-source cooperation, they expect it to be unstable with a short half-life. Since there's no obvious way to enforce an allocation policy for programmer time over the internet, this model leads straight to a prediction that the commons will break up, with various bits of software being taken closed-source and a rapidly decreasing amount of work being fed back into the communal pool. *kinda w/ branching*

In fact, it is empirically clear that the trend is opposite to this. The breadth and volume of open-source development (as measured by, for example, submissions per day at Metalab or announcements per day at freshmeat.net) is steadily increasing. Clearly there is some critical way in which the "Tragedy of the Commons" model fails to capture what is actually going on.

Part of the answer certainly lies in the fact that using software does not decrease its value. Indeed, widespread use of open-source software tends to *increase* its value, as users fold in their own fixes and features (code patches). In this inverse commons, the grass grows taller when it's grazed on. *but com-ersly*

Another part of the answer lies in the fact that the putative market value of small patches to a common source base is hard to capture. Supposing I write a fix for an irritating bug, and suppose many people realize the fix has money value; how do I collect from all those people? Conventional payment systems have high enough overheads to make this a real problem for the sorts of micropayments that would usually be appropriate. *bad UI*  
*No drastic change*  
*Yeah stupid*

It may be more to the point that this value is not merely hard to capture, in the general case it's hard to even *assign*. As a thought experiment let us suppose that the Internet came equipped with the theoretically ideal micropayment system – secure, universally accessible, zero-overhead. Now let's say you have written a patch labeled "Miscellaneous Fixes to the Linux Kernel". How do you know what price to ask? How would a potential buyer, not having seen the patch yet, know what is reasonable to pay for it?



What we have here is almost like a funhouse-mirror image of F.A. Hayek's 'calculation problem' – it would take a superbeing, both able to evaluate the functional worth of patches and trusted to set prices accordingly, to lubricate trade.

Unfortunately, there's a serious superbeing shortage, so patch author J. Random Hacker is left with two choices: sit on the patch, or throw it into the pool for free. The first choice gains nothing. The second choice may gain nothing, or it may encourage reciprocal giving from others that will address some of J. Random's problems in the future. The second choice, apparently altruistic, is actually optimally selfish in a game-theoretic sense.

In analyzing this kind of cooperation, it is important to note that while there is a free-rider problem (work may be underprovided in the absence of money or money-equivalent compensation) it is not one that scales with the number of end-users. The complexity and communications overhead of an open-source project is almost entirely a function of the number of developers involved; having more end-users who never look at source costs effectively nothing. It may increase the rate of silly questions appearing on the project mailing lists, but this is relatively easily forestalled by maintaining a Frequently Asked Questions list and blithely ignoring questioners who have obviously not read it (and in fact both these practices are typical).

The real free-rider problems in open-source software are more a function of friction costs in submitting patches than anything else. A potential contributor with little stake in the cultural reputation game (see 16 ()) may, in the absence of money compensation, think "It's not worth submitting this fix because I'll have to clean up the patch, write a ChangeLog entry, and sign the FSF assignment papers...". It's for this reason that the number of contributors (and, at second order, the success of) projects is strongly and inversely correlated with the number of hoops each project makes a user go through to contribute. Such friction costs may be political as well as mechanical. Together they may explain why the loose, amorphous Linux culture has attracted orders of magnitude more cooperative energy than the more tightly organized and centralized BSD efforts and why the Free Software Foundation has receded in relative importance as Linux has risen.

This is all good as far as it goes. But it's an after-the-fact explanation of what J. Random Hacker does with his patch after he has it. The other half we need is an economic explanation of how JRH was able to write that patch in the first place, rather than having to work on a closed-source program that might have returned him sale value. What business models create niches in which open-source development can flourish?

## 6 Reasons for Closing Source

Before taxonomizing open-source business models, we should deal with exclusion payoffs in general. What exactly are we protecting when we close source?

Let's say you hire someone to write to order (say) a specialized accounting package for your business. That problem won't be solved any better if the sources are closed rather than open; the only rational reasons you might want them to be closed is if you want to sell the package to other people, or deny its use to competitors.

The obvious answer is that you're protecting sale value, but for the 95% of software written for internal use this doesn't apply. So what other gains are there in being closed?

That second case (protecting competitive advantage) bears a bit of examination. Suppose you open-source that accounting package. It becomes popular and benefits from improvements made by the community. Now, your competitor also starts to use it. The competitor gets the benefit without paying the development

right no value of patches - but is value in higher level  
UI stuff  
which is  
Proprietary  
-Ting  
-Palm  
Carter  
Maemo  
CS ga



cost and cuts into your business. Is this an argument against open-sourcing?

Maybe – and maybe not. The real question is whether your gain from spreading the development load exceeds your loss due to increased competition from the free rider. Many people tend to reason poorly about this tradeoff through (a) ignoring the functional advantage of recruiting more development help. (b) not treating the development costs as sunk, and By hypothesis, you had to pay th development costs anyway, so counting them as a cost of open-sourcing (if you choose to do) is mistaken.

There are other reasons for closing source that are outright irrational. You might, for example, be laboring under the delusion that closing the sources will make your business systems more secure against crackers and intruders. If so, I recommend therapeutic conversation with a cryptographer immediately. The really professional paranoids know better than to trust the security of closed-source programs, because they've learned through hard experience not to. Security is an aspect of reliability; only algorithms and implementations that have been thoroughly peer-reviewed can possibly be trusted to be secure.

## 7 Use-Value Funding Models

A key fact that the distinction between use and sale value allows us to notice is that only *sale value* is threatened by the shift from closed to open source; use value is not.

If use value rather than sale value is really the major driver of software development, and (as was argued in 16 ()) open-source development is really more effective and efficient than closed, then we should expect to find circumstances in which expected use value alone sustainably funds open-source development.

And in fact it is not difficult to identify at least two important models in which full-time developer salaries for open-source projects are funded strictly out of use value.

### 7.1 The Apache case: cost-sharing

Let's say you work for a firm that has a business-critical requirement for a high-volume, high-reliability web server. Maybe it's for electronic commerce, maybe you're a high-visibility media outlet selling advertising, maybe you're a portal site. You need 24/7 uptime, you need speed, and you need customizability.

How are you going to get these things? There are three basic strategies you can pursue:

*Buy a proprietary webserver.* In this case, you are betting that the vendor's agenda matches yours and that the vendor has the technical competence to implement properly. Even assuming both these things to be true, the product is likely to come up short in customizability; you will only be able to modify it through the hooks the vendor has chosen to provide. This proprietary-webserver path is not a popular one.

*Roll your own.* Building your own webserver is not an option to dismiss instantly; webserver are not very complex, certainly less so than browsers, and a specialized one can be very lean and mean. Going this path, you can get the exact features and customizability you want, though you'll pay for it in development time. Your firm may also find it has a problem when you retire or leave.

*Join the Apache group.* The Apache server was built by an Internet-connected group of webmasters who realized that it was smarter to pool their efforts into improving one code base than to run a large number of parallel development efforts. By doing this they were able to capture both most of the advantages of roll-your-own and the powerful debugging effect of massively-parallel peer review.

but he patches + maintain

plus I would say most is

not task driven VI

-but

low level bits

-OpenOffice.org

-open

tax

Software

✓ This is infrastructure here applicable to a lot of people not very big centric

patches valuable



The advantage of the Apache choice is very strong. Just how strong, we may judge from the monthly Netcraft survey, which has shown Apache steadily gaining market share against all proprietary web servers since its inception. As of June 1999, Apache and its derivatives have 61% market share <<http://www.netcraft.com/survey/>> – with no legal owner, no promotion, and no contracted service organization behind them at all.

The Apache story generalizes to a model in which software users find it to their advantage to fund open-source development because doing so gets them a better product than they could otherwise have, at lower cost.

## 7.2 The Cisco case: risk-spreading

Some years ago, two programmers at Cisco (the networking-equipment manufacturer) got assigned the job of writing a distributed print-spooling system for use on Cisco's corporate network. This was quite a challenge. Besides supporting the ability for arbitrary user A to print at arbitrary printer B (which might be in the next room or a thousand miles away), the system had to make sure that in the event of a paper-out or toner-low condition the job would get rerouted to an alternate printer near the target. The system also needed to be able to report such problems to a printer administrator.

The duo came up with a clever set of modifications to the standard Unix print-spooler software, plus some wrapper scripts, that did the job. Then they realized that they, and Cisco, had a problem.

The problem was that neither of them was likely to be at Cisco forever. Eventually, both programmers would be gone, and the software would be unmaintained and begin to rot (that is, to gradually fall out of sync with real-world conditions). No developer likes to see this happen to his or her work, and the intrepid duo felt Cisco had paid for a solution under the not unreasonable expectation that it would outlast their own jobs there.

Accordingly, they went to their manager and urged him to authorize the release of the print spooler software as open source. Their argument was that Cisco would have no sale value to lose, and much else to gain. By encouraging the growth of a community of users and co-developers spread across many corporations, Cisco could effectively hedge against the loss of the software's original developers.

The Cisco story generalizes to a model in which open source functions not so much to lower costs as to spread risk. All parties find that the openness of the source, and the presence of a collaborative community funded by multiple independent revenue streams, provides a fail-safe that is itself economically valuable – sufficiently valuable to drive funding for it.

## 8 Why Sale Value is Problematic

again infrastructure  
and many different needs

Open source makes it rather difficult to capture direct sale value from software. The difficulty is not technical; source code is no more nor less copyable than binaries, and the enforcement of copyright and license laws permitting capture of sale value would not by necessity be any more difficult for open-source products than it is for closed.

The difficulty lies rather with the nature of the social contract that supports open-source development. For three mutually reinforcing reasons, the major open-source licenses prohibit most of the sort of restrictions on use, redistribution and modification that would facilitate direct-sale revenue capture. To understand

Some – controversial

plus  
easy  
copying-  
see other  
article



these reasons. we must examine the social context within which the licenses evolved; the Internet *hacker* <<http://www.tuxedo.org/~esr/faqs/hacker-howto.html>> culture.

Despite myths about the hacker culture still (in 1999) widely believed outside it, none of these reasons has to do with hostility to the market. While a minority of hackers does indeed remain hostile to the profit motive, the general willingness of the community to cooperate with for-profit Linux packagers like Red Hat, SUSE, and Caldera demonstrates that most hackers will happily work with the corporate world when it serves their ends. The real reasons hackers frown on direct-revenue-capture licenses are more subtle and interesting.

One reason has to do with symmetry. While most open-source developers do not intrinsically object to others profiting from their gifts, most also demand that no party (with the possible exception of the originator of a piece of code) be in a *privileged* position to extract profits. J. Random Hacker is willing for Fubarco to profit by selling his software or patches, but only so long as JRH himself could also potentially do so.

Another has to do with unintended consequences. Hackers have observed that licenses that include restrictions on and fees for 'commercial' use or sale (the most common form of attempt to recapture direct sale value, and not at first blush an unreasonable one) have serious chilling effects. A specific one is to cast a legal shadow on activities like redistribution in inexpensive CD-ROM anthologies, which we would ideally like to encourage. More generally, restrictions on use/sale/modification/distribution (and other complications in licensing) exact an overhead for conformance tracking and (as the number of packages people deal with rises) a combinatorial explosion of perceived uncertainty and potential legal risk. This outcome is considered harmful, and there is therefore strong social pressure to keep licenses simple and free of restrictions.

The final and most critical reason has to do with preserving the peer-review, gift-culture dynamic described in 16 (). License restrictions designed to protect intellectual property or capture direct sale value often have the effect of making it legally impossible to fork the project (this is the case, for example, with Sun's so-called "Community Source" licenses for Jini and Java). While forking is frowned upon and considered a last resort (for reasons discussed at length in 16 ()), it's considered critically important that that last resort be present in case of maintainer incompetence or defection (e.g. to a more closed license).

The hacker community has some give on the symmetry reason; thus, it tolerates licenses like Netscape's NPL that give some profit privileges to the originators of the code (specifically in the NPL case, the exclusive right to use the open-source Mozilla code in derivative products including closed source). It has less give on the unintended-consequences reason, and none on preserving the option to fork (which is why Sun's Java and Jini 'Community License' schemes have been largely rejected by the community).

These reasons explain the clauses of the Open Source Definition, which was written to express the consensus of the hacker community about the critical features of the standard licenses (the GPL, the BSD license, the MIT License, and the Artistic License). These clauses have the effect (though not the intention) of making direct sale value very hard to capture.

## 9 Indirect Sale-Value Models

Nevertheless, there are ways to make markets in software-related services that capture something like indirect sale value. There are five known and two speculative models of this kind (more may be developed in the future).

Stallman

GPL kinda

I do not  
see commercial  
but most  
want to be  
paid -  
now I see  
how it  
distorts  
model

or must open changes - now controversial



## 9.1 Loss-Leader/Market Positioner

In this model, you use open-source software to create or maintain a market position for proprietary software that generates a direct revenue stream. In the most common variant, open-source client software enables sales of server software, or subscription/advertising revenue associated with a portal site.

Netscape Communications, Inc. was pursuing this strategy when it open-sourced the Mozilla browser in early 1998. The browser side of their business was at 13% of revenues and dropping when Microsoft first shipped Internet Explorer. Intensive marketing of IE (and shady bundling practices that would later become the central issue of an antitrust lawsuit) quickly ate into Netscape's browser market share, creating concern that Microsoft intended to monopolize the browser market and then use de-facto control of HTML to drive Netscape out of the server market.

By open-sourcing the still-widely-popular Netscape browser, Netscape effectively denied Microsoft the possibility of a browser monopoly. They expected that open-source collaboration would accelerate the development and debugging of the browser, and hoped that Microsoft's IE would be reduced to playing catch-up and prevented from exclusively defining HTML.

This strategy worked. In November 1998 Netscape actually began to regain business-market share from IE. By the time Netscape was acquired by AOL in early 1999, the competitive advantage of keeping Mozilla in play was sufficiently clear that one of AOL's first public commitments was to continue supporting the Mozilla project, even though it was still in alpha stage.

## 9.2 Widget Frosting

This model is for hardware manufacturers (hardware, in this context, includes anything from Ethernet or other peripheral boards all the way up to entire computer systems). Market pressures have forced hardware companies to write and maintain software (from device drivers through configuration tools all the way up to the level of entire operating systems), but the software itself is not a profit center. It's an overhead – often a substantial one.

In this situation, opening source is a no-brainer. There's no revenue stream to lose, so there's no downside. What the vendor gains is a dramatically larger developer pool, more rapid and flexible response to customer needs, and better reliability through peer review. It gets ports to other environments for free. It probably also gains increased customer loyalty as its customers' technical staffs put increasing amounts of time into the code to do the customizations they require.

There are a couple of vendor objections commonly raised specifically to open-sourcing hardware drivers. Rather than mix them with discussion of more general issues here, I have written an 17 (appendix) specifically on this topic.

The 'future-proofing' effect of open source is particularly strong with respect to widget frosting. Hardware products have a finite production and support lifetime; after that, the customers are on their own. But if they have access to driver source and can patch them as needed, they're more likely to be happier repeat customers of the same company.

A very dramatic example of adopting the widget frosting model was Apple Computer's decision in mid-March 1999 to open-source "Darwin", the core of their MacOSX server operating system.



### 9.3 Give Away the Recipe, Open A Restaurant

In this model, one open-sources software to create a market position not for closed software (as in the Loss-Leader/Market-Positioner case) but for services.

(I used to call this 'Give Away the Razor, Sell Razor Blades', The coupling is not really as tight as the razor/razor-blade analogy implies.)

This is what Red Hat and other Linux distributors do. What they are actually selling is not the software, the bits itself, but the value added by assembling and testing a running operating system that is warranted (if only implicitly) to be merchantable and to be plug-compatible with other operating systems carrying the same brand. Other elements of their value proposition include free installation support and the provision of options for continuing support contracts.

The market-building effect of open source can be extremely powerful, especially for companies which are inevitably in a service position to begin with. One very instructive recent case is Digital Creations, a website-design house started up in 1998 that specializes in complex database and transaction sites. Their major tool, the intellectual-property crown jewels of the company, is an object publisher that has been through several names and incarnations but is now called Zope.

When the Digital Creations people went looking for venture capital, the VC they brought in carefully evaluated their prospective market niche, their people, and their tools. He then recommended that Digital Creations take Zope to open source.

By traditional software-industry standards, this looks like an absolutely crazy move. Conventional business-school wisdom has it that core intellectual property like Zope is a company's crown jewels, never under any circumstances to be given away. But the VC had two related insights. One is that Zope's true core asset is actually the brains and skills of its people. The second is that Zope is likely to generate more value as a market-builder than as a secret tool.

never heard of it

To see this, compare two scenarios. In the conventional one, Zope remains Digital Creations's secret weapon. Let's stipulate that it's a very effective one. As a result, the firm will be able to deliver superior quality on short schedules – *but nobody knows that*. It will be easy to satisfy customers, but harder to build a customer base to begin with.

The VC, instead, saw that open-sourcing Zope could be critical advertising for Digital Creations's *real* asset – its people. He expected that customers evaluating Zope would consider it more efficient to hire the experts than to develop in-house Zope expertise.

One of the Zope principals has since confirmed very publicly that their open-source strategy has "opened many doors we wouldn't have got in otherwise". Potential customers do indeed respond to the logic of the situation – and Digital Creations, accordingly, is prospering.

Another up-to-the-minute example is *e-smith, inc.* <<http://www.e-smith.net/>>. This company sells support contracts for turnkey Internet server software that is open-source, a customized Linux. One of the principals, describing the spread of free downloads of e-smith's software, says <<http://www.globetechnology.com/gam/News/19990625/BAND.html>> "Most companies would consider that software piracy; we consider it free marketing".

the people w/ the \$ pay for it  
^  
bizs



### 9.4 Accessorizing

In this model, you sell accessories for open-source software. At the low end, mugs and T-shirts; at the high end, professionally-edited and produced documentation.

O'Reilly Associates, publishers of many excellent references volumes on open-source software, is a good example of an accessorizing company. O'Reilly actually hires and supports well-known open-source hackers (such as Larry Wall and Brian Behlendorf) as a way of building its reputation in its chosen market.

### 9.5 Free the Future, Sell the Present

In this model, you release software in binaries and source with a closed license, but one that includes an expiration date on the closure provisions. For example, you might write a license that permits free redistribution, forbids commercial use without fee, and guarantees that the software come under GPL terms a year after release or if the vendor folds.

Under this model, customers can ensure that the product is customizable to their needs, because they have the source. The product is future-proofed – the license guarantees that an open source community can take over the product if the original company dies.

Because the sale price and volume are based on these customer expectations, the original company should enjoy enhanced revenues from its product versus releasing it with an exclusively closed source license. Furthermore, as older code is GPLed, it will get serious peer review, bug fixes, and minor features, which removes some of the 75% maintenance burden on the originator. *roll into modern software*

This model has been successfully pursued by Aladdin Enterprises, makers of the popular Ghostscript program (a PostScript interpreter that can translate to the native languages of many printers).

The main drawback of this model is that the closure provisions tend to inhibit peer review and participation early in the product cycle, precisely when they are needed most.

### 9.6 Free the Software, Sell the Brand

This is a speculative business model. You open-source a software technology, retain a test suite or set of compatibility criteria, then sell users a brand certifying that their implementation of the technology is compatible with all others wearing the brand. *Firefox*

(This is how Sun Microsystems ought to be handling Java and Jini.)

*Ubuntu no software free  
- well FF has*

*Open NFC stacks*

### 9.7 Free the Software, Sell the Content

This is another speculative business model. Imagine something like a stock-ticker subscription service. The value is neither in the client software nor the server but in providing objectively reliable information. So you open-source all the software and sell subscriptions to the content. As hackers port the client to new platforms and enhance it in various ways, your market automatically expands.

(This is why AOL ought to open-source its client software.)

*Mozilla - Google Search deal*

*Me - Hire full time / consulting*

*I think some do this  
but can't think → FF?  
- perhaps hosted - wordpress  
How are they doing?*



## 10 When To Be Open, When To Be Closed

Having reviewed business models that support open-source software development, we can now approach the general question of when it makes economic sense to be open-source and when to be closed-source. First, we must be clear what the payoffs are from each strategy.

### 10.1 What Are the Payoffs?

The closed-source approach allows you to collect rent from your secret bits; on the other hand, it forecloses the possibility of truly independent peer review. The open-source approach sets up conditions for independent peer review, but you don't get rent from your secret bits. *and makes you do all the maintenance yourself*

The payoff from having secret bits is well understood; traditionally, software business models have been constructed around it. Until recently, the payoff from independent peer review was not well understood. The Linux operating system, however, drives home a lesson that we should probably have learned years ago from the history of the Internet's core software and other branches of engineering – that open-source peer review is the only scalable method for achieving high reliability and quality. *true*

In a competitive market, therefore, customers seeking high reliability and quality will reward software producers who go open-source and discover how to maintain a revenue stream in the service, value-add, and ancillary markets associated with software. This phenomenon is what's behind the astonishing success of Linux, which came from nowhere in 1996 to over 17% in the business server market by the end of 1998 and seems on track to dominate that market within two years (in early 1999 IDC projected that Linux would grow faster than all other operating systems combined through 2003). *ha*

An almost equally important payoff of open source is its utility as a way to propagate open standards and build markets around them. The dramatic growth of the Internet owes much to the fact that nobody owns TCP/IP; nobody has a proprietary lock on the core Internet protocols. *well in devices*

The network effects behind TCP/IP's and Linux's success are fairly clear and reduce ultimately to issues of trust and symmetry – potential parties to a shared infrastructure can rationally trust it more if they can see how it works all the way down, and will prefer an infrastructure in which all parties have symmetrical rights to one in which a single party is in a privileged position to extract rents or exert control.

It is not, however, actually necessary to assume network effects in order for symmetry issues to be important to software consumers. No software consumer will rationally choose to lock itself into a supplier-controlled monopoly by becoming dependent on closed source if any open-source alternative of acceptable quality is available. This argument gains force as the software becomes more critical to the software consumer's business – the more vital it is, the less the consumer can tolerate having it controlled by an outside party.

Finally, an important customer payoff of open-source software related to the trust issue is that it's future-proof. If sources are open, the customer has some recourse if the vendor goes belly-up. This may be particularly important for widget frosting, since hardware tends to have short life cycles, but the effect is more general and translates into increased value for open-source software.

### 10.2 How Do They Interact?

When the rent from secret bits is higher than the return from open source, it makes economic sense to be closed-source. When the return from open source is higher than the rent from secret bits, it makes sense to go open source.



In itself, this is a trivial observation. It becomes nontrivial when we notice that the payoff from open source is harder to measure and predict than the rent from secret bits – and that said payoff is grossly underestimated much more often than it is overestimated. Indeed, until the mainstream business world began to rethink its premises following the Mozilla source release in early 1998, the open-source payoff was incorrectly but very generally assumed to be zero.

So how can we evaluate the payoff from open source? It's a difficult question in general, but we can approach it as we would any other predictive problem. We can start from observed cases where the open-source approach has succeeded or failed. We can try to generalize to a model which gives at least a qualitative feel for the contexts in which open source is a net win for the investor or business trying to maximize returns. We can then go back to the data and try to refine the model.

From the analysis presented in 16 (), we can expect that open source has a high payoff where (a) reliability/stability/scalability are critical, and (b) correctness of design and implementation is not readily verified by means other than independent peer review. (The second criterion is met in practice by most non-trivial programs.)

*but is that payoff to you or society?*

A consumer's rational desire to avoid being locked into a monopoly supplier will increase its interest in open source (and, hence, the competitive-market value for suppliers of going open) as the software becomes more critical to that consumer. Thus, another criterion (c) pushes towards open source when the software is a business-critical capital good (as, for example, in many corporate MIS departments).

As for application area, we observed above that open-source infrastructure creates trust and symmetry effects that, over time, will tend to attract more customers and to outcompete closed-source infrastructure; and it is often better to have a smaller piece of such a rapidly-expanding market than a bigger piece of a closed and stagnant one. Accordingly, for infrastructure software, an open-source play for ubiquity is quite likely to have a higher long-term payoff than a closed-source play for rent from intellectual property.

In fact, the ability of potential customers to reason about the future consequences of vendor strategies and their reluctance to accept a supplier monopoly implies a stronger constraint; without already having overwhelming market power, you can choose either an open-source ubiquity play or a direct-revenue-from-closed-source play – but not both. (Analogues of this principle are visible elsewhere, e.g. in electronics markets where customers often refuse to buy sole-source designs.) The case can be put less negatively:

We may sum up this logic by observing that open source seems to be most successful in generating greater returns than closed source in software that (d) establishes or enables a common computing and communications infrastructure.

Finally, we may note that purveyors of unique or just highly differentiated services have more incentive to fear copying of their methods by competitors than do vendors of services for which the critical algorithms and knowledge bases are well understood. Accordingly, open source is more likely to dominate when (e) key methods (or functional equivalents) are part of common engineering knowledge.

The Internet core software, Apache, and Linux's implementation of the ANSI-standard Unix API are prime exemplars of all five criteria. The path towards open source in the evolution of such markets are well-illustrated by the reconvergence of data networking on TCP/IP in the mid-1990s following fifteen years of failed empire-building attempts with closed protocols such as DECNET, XNS, IPX, and the like.

On the other hand, open source seems to make the least sense for companies that have unique possession of a value-generating software technology (strongly fulfilling criterion (e)) which is (a) relatively insensitive to failure, which can (b) readily be verified by means other than independent peer review, which is not

*Was in MIT dining Contract*

*big stability slowly built on SW you control no script replace how is MIT ISAT in this would be interesting to see*

*and also everyone needs it*



(c) business-critical, and which would not have its value substantially increased by (d) network effects or ubiquity.

As an example of this extreme case, in early 1999 I was asked "Should we go open source?" by a company that writes software to calculate cutting patterns for sawmills that want to extract the maximum yardage of planks from logs. My conclusion was "No." The only criterion this comes even close to fulfilling is (c); but at a pinch, an experienced operator could generate cut patterns by hand.

An important point is that where a particular product or technology sits on these scales may change over time, as we'll see in the following case study.

In summary, the following discriminators push towards open source:

- (a)  
reliability/stability/scalability are critical
- (b)  
correctness of design and implementation cannot readily be verified by means other than independent peer review
- (c)  
the software is critical to the user's control of his/her business
- (d)  
the software establishes or enables a common computing and communications infrastructure
- (e)  
key methods (or functional equivalents of them) are part of common engineering knowledge.

### 10.3 Doom: A Case Study

The history of id software's best-selling game Doom illustrates ways in which market pressure and product evolution can critically change the payoff magnitudes for closed vs. open source.

When Doom was first released in late 1993, its first-person, real-time animation made it utterly unique (the antithesis of criterion (e)). Not only was the visual impact of the technique stunning, but for many months nobody could figure out how it had been achieved on the underpowered microprocessors of that time. These secret bits were worth some very serious rent. In addition, the potential payoff from open source was low. As a solo game, the software (a) incurred tolerably low costs on failure, (b) not tremendously hard to verify, (c) not business-critical for any consumer, (d) did not benefit from network effects. It was economically rational for Doom to be closed source.

However, the market around Doom did not stand still. Would-be competitors invented functional equivalents of its animation techniques, and other "first-person shooter" games like Duke Nukem began to appear. As these games ate into Doom's market share the value of the rent from secret bits went down.

On the other hand, efforts to expand that share brought on new technical challenges – better reliability, more game features, a larger user base, and multiple platforms. With the advent of multiplayer 'deathmatch' play and Doom gaming services, the market began to display substantial network effects. All this was demanding programmer-hours that id would have preferred to spend on the next game.

*Games bad for open source for all the reasons*



All of these trends raised the payoff from opening the source. At some point the payoff curves crossed over and it became economically rational for id to open up the Doom source and shift to making money in secondary markets such as game-scenario anthologies. And sometime after this point, it actually happened. The full source for Doom was released in late 1997.

5 years after release

#### 10.4 Knowing When To Let Go

Doom makes an interesting case study because it is neither an operating system nor communication/networking software; it is thus far removed from the usual and obvious examples of open-source success. Indeed, Doom's life cycle, complete with crossover point, may be coming to typify that of applications software in today's code ecology – one in which communications and distributed computation both create serious robustness/reliability/scalability problems only addressable by peer review, and frequently cross boundaries both between technical environments and between competing actors (with all the trust and symmetry issues that implies).

there is the value

Doom evolved from solo to deathmatch play. Increasingly, the network effect is the computation. Similar trends are visible even in the heaviest business applications, such as ERPs, as businesses network ever more intensively with suppliers and customers – and, of course, they are implicit in the whole architecture of the World Wide Web. It follows that almost everywhere, the open-source payoff is steadily rising.

If present trends continue, the central challenge of software technology and product management in the next century will be knowing when to let go – when to allow closed code to pass into the open-source infrastructure in order to exploit the peer-review effect and capture higher returns in service and other secondary markets.

There are obvious revenue incentives not to miss the crossover point too far in either direction. Beyond that, there's a serious opportunity risk in waiting too long – you could get scooped by a competitor going open-source in the same market niche.

The reason this is a serious issue is that both the pool of users and the pool of talent available to be recruited into open-source cooperation for any given product category is limited, and recruitment tends to stick. If two producers are the first and second to open-source competing code of roughly equal function, the first is likely to attract the most users and the most and best-motivated co-developers; the second will have to take leavings. Recruitment tends to stick, as users gain familiarity and developers sink time investments in the code itself.

Especially now if he updates paper

### 11 The Business Ecology of Open Source

The open-source community has organized itself in a way that tends to amplify the productivity effects of open source. In the Linux world, in particular, it's an economically significant fact that there are multiple competing Linux distributors which form a tier separate from the developers.

Developers write code, and make the code available over the Internet. Each distributor selects some subset of the available code, integrates and packages and brands it, and sells it to customers. Users choose among distributions, and may supplement a distribution by downloading code directly from developer sites.

The effect of this tier separation is to create a very fluid internal market for improvements. Developers compete with each other, for the attention of distributors and users, on the quality of their software. Distributors compete for user dollars on the appropriateness of their selection policies, and on the value they can add to the software.

how Ubuntu - threw real \$ behind it



A first-order effect of this internal market structure is that no node in the net is indispensable. Developers can drop out; even if their portion of the code base is not picked up directly by some other developer, the competition for attention will tend to rapidly generate functional alternatives. Distributors can fail without damaging or compromising the common open-source code base. The ecology as a whole has a more rapid response to market demands, and more capability to resist shocks and regenerate itself, than any monolithic vendor of a closed-source operating system can possibly muster. *well maintainers have rights*

Another important effect is to lower overhead and increase efficiency through specialization. Developers don't experience the pressures that routinely compromise conventional closed projects and turn them into tar-pits – no lists of pointless and distracting check-list features from Marketing, no management mandates to use inappropriate and outdated languages or development environments, no requirement to re-invent wheels in a new and incompatible way in the name of product differentiation or intellectual-property protection, and (most importantly) *no deadlines*. No rushing a 1.0 out the door before it's done right – which (as DeMarco and Lister observed in their discussion of the 'wake me when it's over' management style in 16 ()) generally conduces not only to higher quality but actually to the most rapid delivery of a truly working result.

Distributors, on the other hand, get to specialize in the things distributors can do most effectively. Freed of the need to fund massive and ongoing software development just to stay competitive, they can concentrate on system integration, packaging, quality assurance, and service. *but Ubuntu moved into coding*

Both distributors and developers are kept honest by the constant feedback from and monitoring by users that is an integral part of the open-source method.

## 12 Coping With Success

*I would have done this class before WP*

The Tragedy of the Commons may not be applicable to open-source development as it happens today, but that doesn't mean there are not any reasons to wonder if the present momentum of the open-source community is sustainable. *Will key players defect from cooperation as the stakes become higher?* *sustains itself*

There are several levels on which this question can be asked. Our 'Comedy of the Commons' counter-story is based on the argument that the value of individual contributions to open source is hard to monetize. But this argument has much less force for firms (like, say, Linux distributors) which already have a revenue stream associated with open source. Their contribution is already being monetized every day. Is their present cooperative role stable?

Examining this question will lead us to some interesting insights about the economics of open-source software in the real world of present time – and about what a true service-industry paradigm implies for the software industry in the future.

On the practical level, applied to the open-source community as it exists now, this question is usually posed in one of two different ways. One: will Linux fragment? Two: conversely, will Linux develop a dominant, quasi-monopolistic player? *Ubuntu – but can never be a monopoly*

The historical analogy many people turn to when suggesting that Linux will fragment is the behavior of the proprietary-Unix vendors in the 1980s. Despite endless talk of open standards, despite numerous alliances and consortia and agreements, proprietary Unix fell apart. The vendors' desire to differentiate their products by adding and modifying OS facilities proved stronger than their interest in growing the total size of the Unix market by maintaining compatibility (and consequently lowering both entry barriers for independent software developers and total cost of ownership for consumers).

This is quite unlikely to happen to Linux, for the simple reason that all the distributors are constrained to



operate from a common base of open source code. It's not really possible for any one of them to maintain differentiation, because the licenses under which Linux code are developed effectively require them to share code with all parties. The moment any distributor develops a feature, all competitors are free to clone it.

Since all parties understand this, nobody even thinks about doing the kinds of maneuvers that fragmented proprietary Unix. Instead, Linux distributors are forced to compete in ways that actually benefit the consumer and the overall market. That is, they must compete on service, support, and their design bets on what interfaces actually conduce to ease installation and use.

The common source base also forecloses the possibility of monopolization. When Linux people worry about this, the name usually muttered is "Red Hat", that of the largest and most successful of the distributors (with somewhere around 90% estimated market share in the U.S.). But it is notable that within days after the May 1999 announcement of Red Hat's long-awaited 6.0 release – before Red Hat's CD-ROMs actually shipped in any quantity – CD-ROM images of the release built from Red Hat's own public FTP site were being advertised by a book publisher and several other CD-ROM distributors at lower prices than Red Hat's expected list.

not any more

Red Hat itself didn't turn a hair at this, because its founders understand very clearly that they do not and cannot own the bits in their product; the social norms of the Linux community forbid that. In a latter-day take on John Gilmore's famous observation that the Internet interprets censorship as damage and routes around it, it has been aptly said that the hacker community responsible for Linux interprets attempts at control as damage and routes around them. For Red Hat to have protested the pre-release cloning of its newest product would have seriously compromised its ability to elicit future cooperation from its developer community.

but it does not care about code

Perhaps more importantly in present time, the software licenses that express these community norms in a binding legal form actively forbid Red Hat from monopolizing the sources of the code their product is based on. The only thing they can sell is a brand/service/support relationship with people who are freely willing to pay for that. This is not a context in which the possibility of a predatory monopoly looms very large.

## 13 Open R&D and the Reinvention of Patronage

There is one other respect in which the infusion of real money into the open-source world is changing it. The community's stars are increasingly finding they can get paid for what they want to do, instead of pursuing open source as a hobby funded by another day job. Corporations like Red Hat, O'Reilly Associates, and VA Linux Systems are building what amount to semi-independent research arms with charters to hire and maintain stables of open-source talent.

Ubuntu

This makes economic sense only if the cost per head of maintaining such a lab can easily be paid out of the expected gains it will enable by growing the firm's market faster. O'Reilly can afford to pay the principal authors of Perl and Apache to do their thing because it expects their efforts will enable it to sell more Perl- and Apache-related books. VA Linux Systems can fund its laboratory branch because improving Linux boosts the use value of the workstations and servers it sells. And Red Hat funds Red Hat Advanced Development Labs to increase the value of its Linux offering and attract more customers.

To strategists from more traditional sectors of the software industry, reared in cultures that regard patent- or trade-secret-protected intellectual property as the corporate crown jewels, this behavior may (despite its market-growing effect) seem inexplicable. Why fund research that every one of your competitors is (by definition) free to appropriate at no cost?

a lot are foundations



There seem to be two controlling reasons. One is that as long as these companies remain dominant players in their market niches, they can expect to capture a proportional lion's share of the returns from the open R&D. Using R&D to buy future profits is hardly a novel idea; what's interesting is the implied calculation that the expected future gains are sufficiently large that these companies can readily tolerate free riders.

While this obvious expected-future-value analysis is a necessary one in a world of hard-nosed capitalists keeping their eyes on ROI, it is not actually the most interesting mode of explanation for star-hiring, because the firms themselves advance a fuzzier one. They will tell you if asked that they are simply doing the right thing by the community they come from. Your humble author is sufficiently well-acquainted with principals at all three of the firms cited above to testify that these protestations cannot be dismissed as humbug. Indeed, I was personally recruited onto the board of VA Linux Systems in late 1998 explicitly so that I would be available to advise them on "the right thing", and have found them far from unwilling to listen when I did so.

An economist is entitled to ask what payoff is involved here. If we accept that talk of doing the right thing is not empty posturing, we should next inquire what self-interest of the firm the "right thing" serves. Nor is the answer, in itself, either surprising or difficult to verify by asking the right questions. As with superficially altruistic behavior in other industries, what these firms actually believe they're buying is goodwill.

Working to earn goodwill, and valuing it as an asset predictive of future market gains, is hardly novel either. What's interesting is the extremely high valuation that the behavior of these firms suggest they put on that goodwill. They're demonstrably willing to hire expensive talent for projects that are not direct revenue generators even during the most capital-hungry phases of the runup to IPO. And, at least so far, the market has actually rewarded this behavior.

The principals of these companies themselves are quite clear about the reasons that goodwill is especially valuable to them. They rely heavily on volunteers among their customer base both for product development and as an informal marketing arm. Their relationship with their customer base is intimate, often relying on personal trust bonds between individuals within and outside the firm.

These observations reinforce a lesson we learned earlier from a different line of reasoning. The relationship between Red Hat/VA/O'Reilly and their customers/developers is not one typical of manufacturing firms. Rather, it carries to an interesting extreme patterns that are characteristic of knowledge-intensive service industries. Looking outside the technology industry, we can see these patterns in (for example) law firms, medical practices, and universities.

We may observe, in fact, that open-source firms hire star hackers for much the same reasons that universities hire star academics. In both cases, the practice is similar in mechanism and effect to the system of aristocratic patronage that funded most fine art until after the Industrial Revolution – a similarity some parties to it are fully aware of.

## 14 Getting There From Here

The market mechanisms for funding (and making a profit from!) open-source development are still evolving rapidly. The business models we've reviewed in this paper probably will not be the last to be invented. Investors are still thinking through the consequences of reinventing the software industry as one with an explicit focus on service rather than closed intellectual property, and will be for some time to come.

This conceptual revolution will have some cost in foregone profits for people investing in the sale-value 5% of the industry; historically, service businesses are not as lucrative as manufacturing businesses (though as any

*I would not have made this argument*

*now Google - for the search hit - there's content*

*- not loaded in  
but everyone  
uses*



doctor or lawyer could tell you, the return to the actual practitioners is often higher). Any foregone profits, however, will be more than matched by benefits on the cost side, as software consumers reap tremendous savings and efficiencies from open-source products. (There's a parallel here to the effects that the displacement of the traditional voice-telephone network by the Internet is having everywhere).

The promise of these savings and efficiencies is creating a market opportunity that entrepreneurs and venture capitalists are now moving in to exploit. As the first draft of this paper was in preparation, Silicon Valley's most prestigious venture-capital firm took a lead stake in the first startup company to specialize in 24/7 Linux technical support. It is generally expected that several Linux- and open-source-related IPOs will be floated before the end of 1999 – and that they will be quite successful. *ha*

Another very interesting development is the beginnings of systematic attempts to make task markets in open-source development. *SourceXchange* <<http://www.sourcexchange.com/process.html>> and *CoSource* <<http://www.cosource.com/>> represent slightly different ways of trying to apply a reverse-auction model to funding open-source development.

The overall trends are clear. We mentioned before IDC's projection that Linux will grow faster than all other operating systems *combined* through 2003. Apache is at 61% market share and rising steadily. Internet usage is exploding, and surveys such as the Internet Operating System Counter show that Linux and other open-source operating systems are already a plurality on Internet hosts and steadily gaining share against closed systems. The need to exploit open-source Internet infrastructure increasingly conditions not merely the design of other software but the business practices and software use/purchase patterns of every corporation there is. These trends, if anything, seem likely to accelerate.

## 15 Conclusion: Life After The Revolution

What will the world of software look like once the open-source transition is complete? *not much of a transition*

For purposes of examining this question, it will be helpful to sort kinds of software by the degree of completeness which the service they offer is describable by open technical standards, which is well correlated with how commoditized the underlying service has become.

This axis corresponds reasonably well to what people are normally thinking when they speak of 'applications' (not at all commoditized, weak or nonexistent open technical standards), 'infrastructure' (commoditized services, strong standards), and 'middleware' (partially commoditized, effective but incomplete technical standards). The paradigm cases today in 1999 would be a word processor (application), a TCP/IP stack (infrastructure), and a database engine (middleware). *so 2/3 good for open*

The payoff analysis we did earlier suggests that infrastructure, applications, and middleware will be transformed in different ways and exhibit different equilibrium mixes of open and closed source. We recall that it also suggested the prevalence of open source in a particular software area would be a function of whether substantial network effects operate there, what the costs of failure are, and to what extent the software is a business-critical capital good.

We can venture some predictions if we apply these heuristics not to individual products but to entire segments of the software market. Here we go:

Infrastructure (the Internet, the Web, operating systems, and the lower levels of communications software that has to cross boundaries between competing parties) will be almost all open source, cooperatively maintained by user consortia and by for-profit distribution/service outfits with a role like that of Red Hat today. ✓



Applications, on the other hand, will have the most tendency to remain closed. There will be circumstances under which the use value of an undisclosed algorithm or technology will be high enough (and the costs associated with unreliability will be low enough, and the risks associated with a supplier monopoly sufficiently tolerable) that consumers will continue to pay for closed software. This is likeliest to remain true in standalone vertical-market applications where network effects are weak. Our lumber-mill example earlier is one such; biometric identification software seems likeliest, of 1999's hot prospects, to be another. right

Middleware (like databases, development tools, or the customized top ends of application protocol stacks) will be more mixed. Whether middleware categories tend to go closed or open seems likely to depend on the cost of failures, with higher cost creating market pressure for more openness. idk well maybe it's

To complete the picture, however, we need to notice that neither 'applications' nor 'middleware' are really stable categories. In 'Knowing When To Let Go' above we saw that individual software technologies seem to go through a natural life cycle from rationally closed to rationally open. The same logic applies in the large. mixed

Applications tend to fall into middleware as standardized techniques develop and portions of the service becomes commoditized. (Databases, for example, became middleware after SQL decoupled front ends from engines.) As middleware services become commoditized, they will in turn tend to fall into the open-source infrastructure – a transition we're seeing in operating systems right now.

In a future that includes competition from open source, we can expect that the eventual destiny of any software technology will be to either die or become part of the open infrastructure itself. While this is hardly happy news for entrepreneurs who would like to collect rent on closed software forever, it does suggest that the software industry as a whole will *remain* entrepreneurial, with new niches constantly opening up at the upper (application) end and a limited lifespan for closed-IP monopolies as their product categories fall into infrastructure.

Finally, of course, this equilibrium will be great for the software consumer driving the process. More and more high-quality software will become permanently available to use and build on instead of being discontinued or locked in somebody's vault. Ceridwen's magic cauldron is, finally, too weak a metaphor – because food is consumed or decays, whereas software sources potentially last forever. The free market, in its widest libertarian sense including *all* un-coerced activity whether trade or gift, can produce perpetually increasing software wealth for everyone.

## 16 Bibliography and Acknowledgements

[CatB] *The Cathedral and the Bazaar* <<http://www.tuxedo.org/~esr/writings/cathedral-bazaar/>>

[HtN] *Homesteading the Noosphere* <<http://www.tuxedo.org/~esr/writings/homesteading/>>

[DL] De Marco and Lister, *Peopleware: Productive Projects and Teams* (New York; Dorset House, 1987; ISBN 0-932633-05-6)

[SH] Shawn Hargreaves has written a good analysis of the applicability of open-source methods to games; *Playing the Open Source Game* <<http://www.talula.demon.co.uk/games.html>>.

Several stimulating discussions with David D. Friedman helped me refine the 'inverse commons' model of open-source cooperation. I am also indebted to Marshall van Alstyne for pointing out the conceptual importance of rivalrous information goods. Ray Ontko of the Indiana Group supplied helpful criticism. A good many people in audiences before whom I gave talks in the year leading up to June 1999 also helped; if



you're one of those, you know who you are.

It's yet another testimony to the open-source model that this paper was substantially improved by email feedback I received within days after release. Lloyd Wood pointed out the importance of open-source software being 'future-proof'. and Doug Dante reminded me of the 'Free the Future' business model. A question from Adam Moorhouse led to the discussion of exclusion payoffs. Lionel Oliviera Gresse gave me a better name for one of the business models. Stephen Turnbull slapped me silly about careless handling of free-rider effects.

## 17 Appendix: Why Closing Drivers Loses A Vendor Money

Manufacturers of peripheral hardware (Ethernet cards, disk controllers, video board and the like) have historically been reluctant to open up. This is changing now, with players like Adaptec and Cyclades beginning to routinely disclose specifications and driver source code for their boards. Nevertheless, there's still resistance out there. In this appendix we attempt to dispel several of the economic misconceptions that sustain it.

*and still are reluctant*

If you are a hardware vendor, you may fear open-sourcing may reveal important things about how your hardware operates that competitors could copy, thus gaining an unfair competitive advantage. Back in the days of three- to five-year product cycles this was a valid argument. Today, the time your competitors' engineers would need to spend copying and understanding the copy is a substantial portion of the product cycle, time they are *not* spending innovating or differentiating their own product. Plagiarism is a trap you *want* your competitors to fall into.)

In any case, these details don't stay hidden for long these days. Hardware drivers are not like operating systems or applications; they're small, easy to disassemble, and easy to clone. Even teenage novice programmers can do this – and frequently do.

There are literally thousands of Linux and FreeBSD programmers out there with both the capability and the motivation to build drivers for a new board. For many classes of device that have relatively simple interfaces and well-known standards (such as disk controllers and network cards) these eager hackers can often prototype a driver as almost rapidly as your own shop could, even without documentation and without disassembling an existing driver.

Even for tricky devices like video cards, there is not much you can do to thwart a clever programmer armed with a disassembler. Costs are low and legal barriers are porous; Linux is an international effort and there is always a jurisdiction in which reverse-engineering will be legal.

For hard evidence that all these claims are true, examine the list of devices supported in the Linux kernel or in the driver subtrees of sites like *Metalab* <<http://metalab.unc.edu/pub/Linux/hardware/!INDEX.html>>, and notice the rate at which new ones are added.

The message? Keeping your driver secret looks attractive in the short run, but is probably bad strategy in the long run (certainly when you're competing with other vendors that are already open). But if you must do it, burn the code into an onboard ROM. Then publish the interface to it. Go open as much as possible, to build your market and demonstrate to potential customers that you believe in your capacity to out-think and out-innovate competitors where it matters.

If you stay closed you will usually get the worst of all worlds – your secrets will get exposed, you won't get free development help, and you won't have wasted your stupider competition's time on cloning. Most importantly, you miss an avenue to widespread early adoption. A large and influential market (the people

*well a lot are Chinese + don't understand  
or where competitors fast – don't strive to differentiate  
think some more*

who manage the servers that run effectively all of the Internet and more than 17% of business data centers) will correctly write your company off as clueless and defensive because you didn't realize these things. Then they'll buy their boards from someone who did.

## 18 History

This is \$Revision: 1.14 \$.

Versions not described here are minor editorial and typo-fix updates.

20 May 1999, version 1.1 – draft.

18 Jun 1999, version 1.2 – first private review version.

24 Jun 1999, version 1.5 – first public release.

24 Jun 1999, version 1.6 – minor update; point at definition of ‘hacker’.

24 Jun 1999, version 1.7 – minor update; clarify criterion (e).

24 Jun 1999, version 1.9 – ‘future-proofing’, the ‘Free the Future’ model, and a new section on exclusion payoffs.

24 Jun 1999, version 1.10 – better name for the ‘Razor Blades’ model.

25 Jun 1999, version 1.13 – corrected 13% claim about Netscape revenues; added better treatment of free-rider effects, corrected list of closed protocols.

25 Jun 1999, version 1.14 – added e-smith, inc.

9 Jul 1999, version 1.15 – new appendix on hardware drivers, and a better explanation of rivalrous goods due to Rich Morin.