The Economics of Open Source Software: A Survey of the Early Literature

AARON SCHIFF*

Department of Economics, The University of Auckland

Abstract

This paper reviews the recent literature on the economics of open source software. Two different sets of issues are addressed. The first looks at the incentives of programmers to participate in open source projects. The second considers the business models used by profit-making firms in the open source industry, and the effects on existing closed source firms. Some possible future research directions are also given.

1 Introduction to open source software

Open source software is a rapidly expanding phenomenon in the computer software industry. A more precise definition of open source software will be given in subsection 1.1, but for now it will suffice to say it is software that is freely distributed and can be freely modified. The open source industry today is large and growing. There are tens of thousands of open source developers worldwide and a good number of firms are trying to make profits out of activities related to open source software. While these firms cannot make profits from the software directly (it is free, after all), they sell complementary products and services, such as service and support, and documentation. More statistics on the size of the open source industry will be given in subsection 1.2.

Open source software has the potential to fundamentally change the economics of the computer software industry. It will affect how traditional software firms structure the incentives that they provide to their programmers, the product markets in which they compete, and the business models and strategies that they use. It will also impact on the career choices of software developers, and provides an additional means for potential programmers to enter the software industry.

Open source thus raises many interesting economic questions. Most of these questions can be grouped into two broad categories. As the majority of the participants in open source projects are unpaid volunteers, the first set of questions concerns the incentives and motivations for these individuals. The second set of questions is related to the business models used by firms in industries related to open source software, and the effects of open source on the existing software industry.

* Mailing address: Department of Economics, University of Auckland, Private Bag 92019, Auckland, New Zealand. E-mail address: a.schiff@auckland.ac.nz The author thanks Adam Warner and a referee for very helpful comments, and the Centre for Research in Network Economics and Communications for financial support.

This paper summarises the current state of the literature on the economics of open source, with special reference to these two groups of questions. But, before turning to the economic issues, it is necessary to understand in more detail the definition and scope of open source software. The next two subsections provide a brief introduction to open source software and discuss the extent of open source development activity.

1.1 What is open source software?

The *source code* of a computer program is the instructions for the program, written in a human-readable format, usually following the syntax of a high-level programming language such as *C* or *Perl*. Instructions in source code cannot be directly executed by a computer and must first be run through a special program called a compiler which produces machine-readable *binary* or *object* code. While source code can easily be read, understood, and modified by a programmer, it is very difficult to understand binary code, and even more difficult to modify it. For this reason, a program only distributed as binary code is called a *closed source* program. A program where the source code is distributed and can be freely modified (without payment of a royalty or fee) by other programmers is, loosely speaking, called *open source* software.¹

Various restrictions may be placed on what can be done with a particular open source program, depending on the license used by the original author(s). Licensing issues are not the focus of this paper; for the details of many different open source licenses, see <http://www.opensource.org>. In general, however, open source software can be modified, extended, adapted, and incorporated into other programs by other programmers, without paying a fee to any previous contributors to the software. The most common restriction placed on such activity is that any future modifications or derivative programs must also be released as open source software.²

1.2 The extent and scope of open source development

A large number, in the order of tens of thousands, of individuals worldwide spend at least some of their time contributing to open source programming.³ A significant number also contribute to related activities such as writing documentation and user support. The website <http://www.freshmeat.net> provides a daily list of new and upgraded open source software, and this list typically contains 20 to 30 programs each day. On 11 February 2002, the website's statistics list 18,540 open source projects and 109,191 registered users.

A wide range of software is available as open source. The most famous example is Linux, a sophisticated computer operating system kernel. Other well-known examples include Apache (a popular and reliable website server), The Gimp (a powerful image manipulation program), and MySQL (a relational database environment). The reader is

¹ *Open source* is still a relatively new term, and is referred to as *free software* in some cases. A definition of free software is provided by the Free Software Foundation: <http://www.fsf.org/philosophy/free-sw.html>. For a more precise definition of open source software, including some subtleties not considered here, see <http://www.opensource.org>.

² The popular General Public License (GPL), among others, imposes the restriction that software must be released as open source if binaries of the software are distributed. That is, modifications made solely for private use do not have to be released as open source.

³ This statistic is a very rough estimate based on browsing open source developer websites, such as <http://www.freshmeat.net> and <http://www.osdn.com>.

recommended to browse <http://www.freshmeat.net> to fully appreciate the scope of open source software that is available and the rapid pace of development.

Until recently, almost all open source developers were people who had a strong interest in computers and did the programming in their spare time, or were academics. However, two new trends have emerged. First, some existing closed source programs have been re-released as open source, such as Netscape's Mozilla web browser.⁴ Second, profit-oriented companies have appeared in the open source community. These companies do not make money from open source software directly, but sell complementary products and services. Two such examples are Red Hat, which distributes Linux and sells service and support packages, and O'Reilly and Associates, which publishes books on open source software. This phenomenon will be discussed in section 2.2, but for now note that such firms will further increase the extent and scope of open source software development activity.

2 The economics of open source software

In this section, the two groups of economic questions relating to open source software mentioned in the introduction are addressed in turn, and the insights from the existing economic literature are explained.

2.1 Incentives to participate

Whilst some programmers are employed to spend some or all of their time working on open source projects by firms like Red Hat or are academics doing research, the majority of contributors to open source projects are unpaid volunteers. What motivates skilled programmers to expend time and effort for no immediate monetary reward? Economic answers to this question have been discussed extensively by Lerner and Tirole (2000).⁵ Explanations put forth by some non-economic commentators include things like altruism, "programming is fun", "for the intellectual challenge", or "to make a political statement". While a few programmers may have political motivations, Lerner and Tirole point out that the altruism hypothesis does not explain why programmers do not focus their attention on more needy causes, and why free riding would be less prevalent than in other industries. Furthermore, they remark that we should be sceptical of the other explanations, since it is unclear why programmers cannot find such challenges in paid employment. We must therefore look for alternative economic explanations of this behaviour.

Basic economic theory tells us that a programmer will participate in an open source project only if it gives a positive net benefit. Lerner and Tirole identify the following immediate benefits and costs. First, the programmer may receive monetary compensation if working for a commercial firm. Second, the programmer may be customising an existing

⁴ See <http://www.mozilla.org>.

⁵ See also section 2.1 of Lerner and Tirole (2001). More generally, the economic literature on 'social culture', such as Fang (2001), could be used as a context for analysing these issues. The economic theory of social culture is beyond the scope of this paper.

program or fixing a bug, which gives a direct benefit.⁶ On the cost side, there is an opportunity cost of time whilst working on the project, which depends on a number of factors including the individual's enjoyment of programming.

Let us compare these immediate costs and benefits between closed and open source projects. First, Lerner and Tirole note that because open source code is freely available, it is often used in universities and schools for teaching, and so may already be familiar to programmers, thus reducing their cost of participating in an open source project. Lerner and Tirole call this the 'alumni effect'. Second, by contributing to an open source project, a programmer may derive some private benefit through bug fixing or customisation that could not be done with a closed source program.

In addition to these immediate costs and benefits, Lerner and Tirole also emphasise the existence of *delayed* payoffs from open source projects. They separate delayed payoffs into two different incentives: the *career concern incentive*, which relates to future job offers or future access to the venture capital market, and the *ego gratification incentive*, which stems from a desire for peer recognition.⁷ Lerner and Tirole group these two incentives together under the heading of the *signalling incentive*, and argue that labour economics, and especially the literature on 'career concerns' such as Holmström (1999), provides useful insights in this regard. From this literature, Lerner and Tirole identify three factors that will increase the signalling incentive: (i) the more visible the performance to the relevant audience, (ii) the higher the impact of effort on performance, and (iii) the more informative the performance about talent.⁸

Lerner and Tirole identify several reasons why the signalling incentive may be greater in open source projects compared to closed source ones. First, performance measurement is better under open source. This is because only the functionality and usage of a closed source program can be observed by outsiders, while under open source the contribution of each individual and the quality of his or her code can be directly inferred. Second, an open source programmer takes full responsibility for his or her project or subproject whereas in a traditional firm environment, an individual's performance depends on that of others. Finally, open source programmers are less likely to have firm-specific human capital, thus making the labour market more fluid.

Lerner and Tirole also argue that these theories can shed light on the type of people that are more likely to contribute to open source projects, and what types of programs are suited to open source. First, sophisticated users are likely to contribute because they can derive direct benefits from customising or bug fixing. Second, individuals with strong signalling incentives, for example someone who lacks formal qualifications but is

⁶ The programmer will get these benefits whether or not the customisation or bug fix is shared with others. However, such improvements are only possible with open source programs because they require access to the source code to implement.

⁷ Eric Raymond, a leader in the open source community, in an influential paper advocating open source development called *The Cathedral and the Bazaar* (2000a), also emphasises the ego gratification incentive: "The 'utility function' Linux hackers are maximizing is not classically economic, but is the intangible of their own ego satisfaction and reputation among other hackers. ... Voluntary cultures that work this way are not actually uncommon; one other in which I have long participated is science fiction fandom, which unlike hackerdom has long explicitly recognized 'egoboo' (ego-boosting, or the enhancement of one's reputation among other fans) as the basic drive behind volunteer activity." This 'egoboo' phenomenon may be one of the motivations for tenured academics to publish papers!

⁸ Note that (i) creates network effects: programmers will want to work on software projects that will attract a large number of other programmers. This may give rise to multiple equilibria, as in Katz and Shapiro (1985).

nevertheless accomplished at programming, may use open source to enter the software industry. With regard to the types of programs best suited to open source, programs whose natural audience is computer programmers, such as operating systems and programming languages, would create the strongest signalling incentives.⁹

An alternative approach to explaining the participation of programmers in open source projects is given by Johnson (2001) who develops a game-theoretic model of programmer behaviour that borrows from the theory of private provision of public goods.¹⁰ This model attempts to explain why purely self-interested programmers may wish to make an improvement to an open source program even when it is known that there are other programmers who may make the same improvement, and even when programmers cannot coordinate their actions.

In Johnson's model there are n user-developers. Each agent i knows of a possible improvement to an open source program and each can develop this enhancement at a private cost c_i . The agents independently choose whether to develop the software, and if it is developed by at least one agent then it is available to all. If the enhancement is developed then all the agents receive their privately known valuations v_i . It is assumed that all agents' costs and valuations are independent and identical draws from a smooth joint distribution $G(c, v)$, with support on a finite rectangle defined by $\{(c, v) : c_L \leq c \leq v_L, v_L \leq v \leq v_H\}$, with $c_L > 0$ and $v_H \geq 0$. Each agent's strategy is a decision to develop or not develop, conditional on their realised cost and valuation, and given their beliefs about what the other agents will do. These beliefs can be conveniently summarised by the probability with which agent i believes that the development will take place if he or she does not develop it, π_i . An agent thus chooses to develop the program if $v_i - c_i > \pi_i v_i$, which implies $v_i / c_i > 1 / (1 - \pi_i)$, that is, if the agent's value-cost ratio is large enough.

Johnson considers symmetric Bayesian Nash equilibria of this game in order to answer a number of questions. First, he analyses what happens when the number of user-developers increases. Obviously, an increase in n makes development more likely simply because it is more likely that there will exist a developer with a high enough value-cost ratio. On the other hand, an increase in n increases the incentive of all agents to free-ride. Therefore, the change in the equilibrium probability of development is ambiguous, and depends on the distribution G . However, without any assumptions about G , Johnson shows that any decline in the development probability arising from an increase in n is bounded and converges to zero as n gets large. This is because, when n is large, the marginal effect of an additional agent on the probability of development is small. Johnson then shows that expected social welfare is increasing in n . This is because each individual is better off in expectation when n increases, due to the fact that the probability that another agent develops the project increases with n .¹¹

Another result Johnson derives is that as n goes to infinity, the equilibrium probability of development is bounded and less than one. Thus, even an infinite number of open source developers may not lead to a desirable enhancement being made. This is because, in the limit, only agents with the highest value-cost ratios will choose to develop

⁹ While there is indeed a bias in open source software towards this type of program, there is also a significant number of 'consumer'-type programs, such as Gnome (a graphical desktop interface), StarOffice (a productivity package including word processor and spreadsheet), as well as numerous games.

¹⁰ See, for example, Chamberlin (1974) and Bliss and Nalebuff (1984).

¹¹ Note, however, that an individual agent and society may be worse off in some states of nature.

the software. Given that the support of value-cost ratios is bounded, the asymptotic equilibrium probability of no development must keep the agent with the highest value-cost ratio indifferent between developing and not developing.¹² On the other hand, it is possible that two or more programmers will simultaneously develop the enhancement. However, Johnson shows that the incentive to free ride puts a bound on such wasteful duplication of effort.

One criticism of Johnson's approach is that it ignores some of the insights provided by Lerner and Tirole's analysis as to the incentives for programmers to participate in open source projects. In particular, recall that Lerner and Tirole argued that a programmer's utility from participating in a project should be increasing in the size of the programmer audience, due to the signalling incentives. Johnson's model, on the other hand, focuses purely on the immediate benefits to the programmer.

2.2 Business models and competitive effects

As mentioned in the introduction, the open source phenomenon has spawned a number of firms that are trying to make profits by selling complementary products and services. Raymond (2000b) identifies seven different business models that such firms can use. These business models are summarised in table 1. Very little theoretical or empirical work has been done on the viability of these business models, competition between such firms, and the effect of firms using these business models on existing closed source software firms.

An economic analysis of the interaction between a closed source monopolist and an open source community is provided by Mustonen (2001). In his model, there is an open source program that is a substitute for a program provided by a profit-maximising monopolist. It is assumed that consumers' valuations of either program are proportional to the amount of development effort that went into the program. Consumers face an 'implementation cost', which is equal for both programs, on top of any price that they must pay to buy the program. There is a large population of programmers who can choose to work for the monopolist at a wage that the monopolist sets, or can choose to belong to the open source community and receive 'complementary income'. In order to increase development effort, the monopolist must hire more programmers, which requires offering a higher wage, while the amount of open source development is determined by the occupational choices of the programmers. The timing of the game is as follows. First, the monopolist sets the wage that it will offer to programmers. Programmers then choose their occupation. Next, consumers evaluate the two programs and the monopolist sets its price level. Finally, the programs are produced and sold.

¹² Johnson notes that this result depends on the support of the value-cost ratio distribution being bounded. If not, development would occur with an arbitrarily high probability as n grew large.

Name	Business model	Example
Loss-leader / Market positioner.	Use open source software to maintain a market position for a related proprietary software product.	Netscape's open source Mozilla web browser and proprietary server software.
Widget frosting.	Sell hardware with open source driver software.	Apple's MacOS X.
Give away the recipe, open a restaurant.	Distribute open source software and sell service and support contracts.	Red Hat.
Accessorising.	Sell accessories for open source software such as documentation.	O'Reilly and Associates.
Free the future, sell the present.	Sell closed source software with a license that makes it open source after a specified time period.	Aladdin's Ghostscript.
Free the software, sell the brand.	Sell other developers a brand that certifies their implementation of your open source technologies is compatible with all others who use the brand.	Sun's StarOffice.
Free the software, sell the content.	Develop an open source product that receives proprietary content that the firm sells.	N/A.

Table 1: Business models related to open source software. Adapted from Raymond (2000b).

Mustonen's model thus considers the effects on a monopolist of constraints in the labour market for programmers and competition from a substitute product that is freely available. Mustonen shows that whether the open source program will be available in the market or not depends on the consumers' implementation costs for programs. In particular, if this cost is sufficiently low, in equilibrium some consumers will choose the open source program and the monopolist will take this into account when choosing its price and will not be able to apply full monopoly power in the product market.

A more heuristic analysis of the effects of the open source movement on closed source firms is provided by Lerner and Tirole (2000). Some advocates of open source, such as Raymond (2000a), argue that it is a superior model of software development in terms of the speed of development and the fixing of bugs, due to the large number and high motivation of open source programmers. In line with this thinking, Lerner and Tirole discuss the extent to which closed source firms can emulate the incentive structure provided by open source. They point out that closed source firms will never be able to replicate the visibility of performance of open source, which will reduce the incentives and motivation of programmers in closed source firms. On the other hand, closed source firms can, and do, list people who have contributed to software,¹³ and promote code-sharing within the firm, to try to restore some of these incentives.

¹³ Lerner and Tirole note that there is a tradeoff from doing this, because it may also induce poaching of talented employees by other firms.

Some closed source software firms encourage their programmers to spend part of their time working on open source projects. Lerner and Tirole identify several reasons why the firms may wish to do this. First, it allows them to have better knowledge about the state of open source development, and hence their competition. Second, it is a way of identifying talented programmers to hire. Third, it may help to provide an intellectually challenging environment for their programmers. Finally, they may want to pre-empt the development of a standard based on a technology owned by a larger rival firm.

Lerner and Tirole also consider the attempts of some closed source firms to re-release their proprietary software under an open source license, such as Netscape's Mozilla web browser. This particular project has so far been relatively unsuccessful in terms of the contributions from open source programmers. Lerner and Tirole suggest that this may be in part due to the fact that a corporation may be unable to credibly commit to keeping all of the source code open and to acknowledging all contributors. Lerner and Tirole then point out that this provides a role for intermediaries such as Collab.Net, which organises open source projects on behalf of other companies, and, as an impartial third party, may be able to provide the needed credibility.

It is clear that much work remains to be done on open source business models and the effects on competition in the software market. The next section briefly summarises some possible avenues of research.

3 Future directions

From the papers reviewed above, it seems that issues surrounding programmer participation in open source projects can be quite well explained by existing economic theory, in particular the theory of private provision of public goods, and various results from labour economics. However, a synthesis of the Lerner and Tirole (2000) labour economics approach and the Johnson (2001) public goods approach may be useful. Future research in the economics of open source software could also focus on business models and competition in the software industry, and on appropriate governance structures for open source projects.

In this regard, one open question considers the positioning of closed source and open source software products. Software users vary in their technical sophistication and requirements. Casual observation suggests that open source software is largely aimed at sophisticated users, while closed source software is often more 'user-friendly'. To some extent this can be explained by the fact that open source programmers seek recognition from their peers, who are sophisticated users. However, it could also be thought of as a product differentiation strategy by the closed source firm(s). The question is then what the optimal degree and determinants of user-friendliness of a closed source program are when faced with open source competition.

Second, as noted above, a possible reaction of closed source firms to the open source phenomenon is to try to emulate some of the open source incentives within the firm. In-depth case studies of modern closed source software development (if possible) would shed light on the extent to which firms are doing this, and the methods that they use. Theoretical work could also give some insight as to the effects of this mimicking behaviour on the open source community, and whether it is likely to be successful for the closed source firms.

Finally, Lerner and Tirole (2001) note that several recent open source projects have been initiated by profit-making firms, either from scratch or by re-releasing existing proprietary software as open source. Therefore, they see one of the key issues as the choice of governance mechanism, to ensure the correct incentives on all sides. As suggested above, this may be facilitated by neutral intermediaries with two-sided reputations. An economic theory of such entities would be very useful, not just for the economics of open source software, but also for understanding a variety of industries.¹⁴

4 References

- Bliss, C. and B. Nalebuff (1984) “Dragon-slaying and ballroom dancing: The private supply of a public good”, *Journal of Public Economics*, 25: 1 – 2.
- Chamberlin, J. (1974) “Provision of Collective Goods as a Function of Group Size”, *American Political Science Review*, 68: 707 – 716.
- Dixit, A. (2001) “On modes of economic governance”, mimeo, Princeton University.
- Fang, H. (2001) “Social culture and economic performance”, *American Economic Review*, 91: 924 – 937.
- Holmström, B. (1999) “Managerial incentive problems: A dynamic perspective”, *Review of Economic Studies*, 66: 169 – 182.
- Johnson, J.P. (2001) “Economics of open source software”, mimeo, <http://opensource.mit.edu/papers/johnsonopensource.pdf>, accessed 4 Feb. 2002.
- Katz, M.L. and C. Shapiro (1985) “Network externalities, competition, and compatibility”, *American Economic Review*, 75: 424 – 440.
- Lerner, J. and J. Tirole (2000) “The simple economics of open source”, *NBER Working Paper*, No. 7600.
- Lerner, J. and J. Tirole (2001) “The open source movement: Key research questions”, *European Economic Review*, 45: 819 – 826.
- Mustonen, M. (2001) “Copyleft – the economics of Linux and other open source software”, mimeo, University of Helsinki.
- Raymond, E.S. (2000a) “The cathedral and the bazaar”, mimeo <http://www.tuxedo.org/~esr/writings/cathedral-bazaar/>, accessed 3 Feb. 2002.
- Raymond, E.S. (2000b) “The magic cauldron”, mimeo, <http://www.tuxedo.org/~esr/writings/magic-cauldron/>, accessed 3 Feb. 2002.

¹⁴ Dixit (2001) provides a starting point for this literature.