"Reverse Anticommons": the Dynamics of Property Rights in Free/Open Source Communities

Suresh Naidu
Maria Alessandra Rossi

June 29, 2004

First draft, preliminary and incomplete

Abstract

We construct a simple stochastic model that incorporates some insights from the New Property Rights approach and use the framework to compare the performance of the F/OSS institutional environment and of more traditional governance structures based on the combination of ownership and markets. The F/OSS licensing structure is regarded as one form of private response to transaction cost-induced frictions in the market for intellectual assets that may take the form of "anticommons" property regimes. Our main result is that, in spite of the incentive effect ownership generates, it might be the case that a different institutional environment that does not rely on such incentives can achieve *on average* better productive outcomes, in circumstances in which transaction-cost induced frictions in the market for intellectual assets tend to give rise to an "anticommons" situation.

1 Introduction

The multifaceted phenomenon that goes under the label of "free" or "open source" software (F/OSS) development has attracted an ever increasing amount of scholarly attention in recent years. The very defining characteristics of F/OSS software, namely a) the free availability of its source code, and b) the nature of the license under which it is distributed, which allows both its free modification and its free redistribution, have raised a number of questions, often related to the underlying issue of whether the phenomenon can be accommodated within the framework of standard economic theory (Lerner and Tirole, 2002).

Scholars in a variety of disciplines have addressed a wide range of questions related to F/OSS, including the nature of individual motivations, the governance structure of single projects, and the likely evolution of competition between open and closed forms of software provision, just to cite a few examples[1]. Some early contributions have considered the peculiarities of the notion of ownership adopted in F/OSS communities, either stressing the anarchistic and non-propertarian nature of F/OSS software production (Eblen, 1999) or emphasizing the Lockean nature of ownership entitlements (Raymond, 1998). However, we are not aware of any systematic attempt to identify through the crafting of a formal model the implications in terms of performance of the peculiar ownership system created by the F/OSS system of licenses.

The main aim of this paper is to compare the performance of the F/OSS institutional environment and of more traditional governance structures based on the combination of ownership and markets[2]. What motivates the paper is the observation that, in our view, F/OSS is suggestive of the fact that, in spite of the undeniable incentive effects the allocation of ownership generates, the costs associated to it may sometimes outweigh its benefits. When this is the case, private agents may react by modifying the very content of ownership by means of contractual arrangements, as it happens with F/OSS licenses. In fact, we argue that the peculiar system of licenses at the heart of F/OSS communities induces a sort of "reverse anticommons", namely - paraphrasing Heller and Eisenberg's metaphor of the "anticommons" - a situation in which *multiple owners each have a right to use and*

[1]For an extensive survey of the literature on F/OSS, see Rossi (2004).

[2]The paper is thus close in spirit to Benkler's (2003) characterization of F/OSS as a "peer production" process. Indeed, Benkler's intuition that *"peer production has a systematic advantage over markets and firms in matching the best available human capital to the best available information inputs"* constitutes a valuable starting point for our modeling effort. Our interest in this paper lies in exploring in greater detail what is it that confers to the F/OSS production model the advantage Benkler has identified and what are the potential shortcomings of such a model.

no one has an effective privilege of exclusion.

In order to perform the comparison between the two institutional arrangements we adapt a simple stochastic model from statistical physics. The model is also inspired by some insights derived from the New Property Rights approach[3], interpreted as a theory of the allocation of intellectual ownership (Pagano and Rossi, 2004). Indeed, the focus of this theory on the incentive role exercised by the allocation of ownership, interpreted in Hart and Moore (1990) as the right to exclude from access to the asset, makes it a natural starting point for a comparison of governance mechanisms based on the combination of ownership and markets with an institutional structure in which agents voluntarily agree to suspend the exercise of the right to exclude with respect to coalition participants. The use of a dynamic model is made necessary by the fact that, in order to assess the limits of conventional ownership-based governance structures, we need to take the NPR approach one step further by introducing positive transaction costs in the market for the exchange of assets and by considering how the frictions in the reallocation of given property rights through time affect productive outcomes.

The model proceeds in three steps: we first obtain as special cases of our model the traditional NPR solution and the case of an "anticommons" property regime and we then perform the comparison between the F/OSS "reverse anticommons" and conventional ownership-based governance structures.

It is worth noting that the results of our model do not depend on particular behavioral assumptions concerning developers. Developers in our model are, in other words, consistent with the description of *homo oeconomicus* common in standard economics textbooks, although alternative assumptions can be easily incorporated into the analysis. This is interesting in light of the fact that the controversy surrounding the interpretation of the F/OSS phenomenon has mainly revolved around the question whether self-regarding preferences can provide a sufficient explanation of developers' behavior. By making "textbook" assumptions concerning the motivation of developers we do not mean to take a particular stance on this issue. Our aim is to construct a framework that allows an assessment of the relevance of the institutional structure in determining productive outcomes, independently of the assumptions made on the nature of agents' preferences.

The paper is organized as follows. In section two, we introduce the no-

[3]With the expression "New Property Rights approach" we refer to a collection of contributions by Sanford Grossman, Oliver Hart and John Moore (among which it should be included: Grossman and Hart 1986; Hart and Moore 1990 and Hart 1995) that is also referred to as the "GHM theory of the firm".

tion of "anticommons" property regimes and we consider some of the factors that may reduce the performance of markets for intellectual assets. Section three describes how the content of ownership rights is contractually modified through the use of F/OSS licenses so as to induce a sort of "reverse anticommons". Section four introduces the model and our main results. Section five derives some implications of our model for the broader issue of the appropriate form of legal protection for software programs. Section six concludes.

2 Markets for Intellectual Assets and "Anticommons" Property Regimes

The granting of intellectual property rights (IPRs) is generally justified on two grounds: because they perform an incentive function and because of their transactional function. The first refers to the fact that IPRs reduce the extent of misappropriation of the information produced and therefore increase the degree of internalization of its benefits by those investing in its production. The second refers to the role played by IPRs as means to solve Arrow's "fundamental paradox" and create the possibility of market exchange of information by allowing the buyer to assess the value of information while protecting the seller from the possibility of loosing its "possession".

Recent scholarship in the intellectual property field has contributed to show that the two functions performed by the IP system are more connected than it would first appear. In particular, it has been argued that the efficiency of private contracting strongly influences incentives to innovate and therefore the ease with which inventors are able to contract over the rearrangement and transfer of intellectual property rights deeply affects the characteristics of optimal IP design, especially in a context of cumulative innovation (Gallini and Scotchmer, 2002). At the same time, the question whether the current IP system facilitates or hampers transactions has also been addressed, spurring a confrontation between "transactional optimists" and "transactional pessimists" (Cooper Dreyfuss, 2003). The first express a more "Coasian" attitude towards exchange to the extent that they stress the importance of a clear definition of rights to enable bargaining. The latter are less sanguine about the likelihood of efficient IPRs exchange and tend to express concern about the possibility that particular features of the IPRs regime may further reduce such likelihood.

The most blatant manifestation of "transactional pessimism" has perhaps been the identification of the risk of a "tragedy of the anticommons",

namely a situation in which *"multiple owners each have a right to exclude others from a scarce resource and no one has an effective privilege of use"*. The hypothesis has been advanced by Michael Heller and Rebecca Eisenberg in a rather influential article appeared in 1998 on *Science* and it then referred to the possibility that the extension of patentability to genes, genetic material and DNA fragments could generate an excessive degree of fragmentation of property rights that would hamper subsequent research because of the substantial transaction costs involved in the attempt to reach a more coherent aggregation of rights. The metaphor of the "anticommons" might sound paradoxical, but is actually consistent with the fundamental nature of patents that grant inventors the right to exclude from access to their intellectual creations, but not the right to use them if they infringe on the rights of other inventors.

The problems with "anticommons" property regimes manifest themselves in more general contexts than intellectual property, namely in any private property regime, and have been recently the object of some research in the law and economics literature (Heller, 1997; Buchanan and Yoon, 2000; Depoorter and Parisi, 2001; Parisi, 2002; Parisi, Shultz and Depoorter, 2003; Shultz, Parisi and Depoorter, 2003). This strand of research has contributed, among other things, to clarify the basic elements of "anticommons" situations, to identify the conditions under which they are likely to produce substantial welfare losses and to propose some mechanisms likely to reduce their adverse consequences. Here we will only recall two basic intuitions. The first is that "anticommons" property regimes result from a lack of conformity between the right to use and the right to exclude that, in turn, generates a negative externality when exclusion rights are exercised simoultaneously and independently by the multiple owners, and generates both a negative externality and hold-up problems when exclusion rights are exercised sequentially. The second is that *"the result of underutilization of joint property increases monotonically in both (a) the extent of fragmentation; and (b) the foregone synergies and complementarities between the property fragments"* (Parisi, Shultz and Depoorter, 2003).

"Anticommons" tragedies are, however, more likely to occur in the intellectual property domain than in the physical property domain, for a variety of reasons that have to do with the obstacles to efficient IPRs exchange pointed out by "transactional pessimists". A first category of impediments relates to search costs. When rights are highly fragmented, as it can easily be the case with rights to composite intellectual creations, the costs incurred in identifying the relevant right-holders are likely to be significant and to be positively correlated to the degree of asset dispersion. The search problem is further complicated by the existence of so-called "submarine" patents, namely patents whose existence is difficult to assess because their applica-

tion is pending.

A second category of impediments arises with respect to negotiation costs and, more specifically, to the issue of valuation. First, there is the problem of asymmetric information. Even though patents disclose some of the information relevant to the determination of assets' value, such information is likely to be at best limited, especially in a context of rapid technological change. Second, as remarked by Heller and Eisenberg (1998), some insights from the behavioral literature[4] suggest that self-serving biases might hinder agreement on valuation. Third, and most importantly, the very nature of intellectual assets affects the ability to strike a bargain. The exact boundaries of rights over intangibles are necessarily ill-defined and generally clearly defined only ex-post, in the context of litigation. Under such circumstances, even agreement about whether a transaction is required at all might be difficult to reach. This is especially apparent in the context of cumulative innovation, when follow-on inventors' ability to commercialize their product depends on obtainment of a license from creators of the invention they build upon. In this context of so-called "blocking patents" the problem of valuation implies a problem of attribution of value to interdependent assets: the extent to which the improvement builds on the previous invention and the relative contribution of the two inventions to total value are, indeed, difficult to determine and bargaining might be affected by strategic behavior (Merges, 1994).

As long as the mentioned transactional difficulties arise *ex-ante*, namely before substantial investments have been sunk, in many technological domains innovators still face the possibility of "inventing around" the assets they have difficulty in acquiring (Shapiro, 2002). However, when the existence of potentially infringed IPRs becomes known only *ex-post*, the "anticommons" problem is at its worst. Inventors unintentionally infringing on previous patents are then exposed to the threat of hold-up by right-holders. As the degree of fragmentation of property rights increases, this threat might become a serious deterrent to innovation as a heavy patent tax burden is imposed on innovators through the threat of multiple hold-ups.

Recent empirical evidence shows that this is indeed a relevant concern in many technological domains characterized by complexity and cumulativeness of innovation. Ziedonis (2000) conducts interviews with representatives from the semiconductor industry and finds evidence of concerns over the costs associated to bargaining with multiple patent owners. Walsh, Arora and Cohen (2003) find similar results in a series of interviews with representatives from the biotech field, both in the private and public sector, although

[4]See Ross and Andersen (1982).

they reject the hypothesis of the existence of a full-blown "tragedy of the anticommons" (but see Rai and Eisenberg 2003 for a rebuttal of this conclusion).

What can we say about the relevance of "anticommons" problems in the software domain? Hints for an answer to this question can be found in the observation that - whatever the specific reason - as a matter of fact there exists at present no market for software components. IP-related transactions, if any, occur almost exclusively ex-post, often in the context of litigation (CSTB, 1999)[5] . Indeed, it is common for developers to build entire programs from scratch, without making recourse to the external market to acquire existing components. Licensing activity thus takes place mainly *ex-post*, exactly when "anticommons" problems are likely to be most thorny.

"Anticommons" property regimes can be looked at as the manifestation of a more general paradox concerning ownership. If one interprets ownership as a vector of rights over all the possible uses of an asset, then it is readily apparent that, in a world of zero transaction costs its role would be insignificant and it is only in presence of positive transaction costs that the important economic function of ownership can be appreciated. This is because, were transaction costs absent (or, put it another way, were contracts complete and fully enforceable) it would, in principle, be possible to contract *ex-ante* on all the possible uses of the asset and ownership would therefore be irrelevant. At the same time, however - and rather paradoxically - it is exactly the existence of positive transaction costs (in the exchange of assets) that reveals the limits of the combination of ownership and markets as a governance mechanism, as the "anticommons" story suggests[6]. These limits are particularly acute in presence of intellectual assets not only because their exchange is exposed to the obstacles pointed out in the previous paragraphs, but also because it is particularly difficult to foresee *ex-ante* all the possible uses of an intellectual asset so as to identify the most suitable owner.

The idea that ownership becomes relevant in presence of positive transaction costs has been emphasized by the New Property Rights approach. This theory has stressed the incentive function performed by the allocation of ownership in a framework characterized by contractual incompleteness and specific investments in human capital. Under such circumstances, agents making investments specific to given assets are exposed to the possibility of hold-up by the other parties to the transaction and the extent to which they are able to appropriate the surplus from their investments depends on their

[5]In terms of our model, this observation justifies very restrictive assumptions as regards to the rate at which the asset part of the control structure mutates under the "market plus property" institutional environment.

[6]Thanks to Ugo Pagano for pointing this out to us.

access to the assets and on the cooperation of the other agents. The reason why ownership matters in this environment is, on one side, that it confers the right to decide over the use of the asset in future uncontracted-for contingencies and, on the other side, that it confers the right to exclude from access to the asset and therefore increases the owner's ex-post bargaining power(Hart and Moore, 1990).

In addition to recognizing the circumstances whereby ownership matters, the New Property Rights approach has also addressed the question of the characterization of the optimal control structure showing that, although the allocation of ownership over non-human assets allows to reach only a second-best solution, the result is in any case superior to that obtainable in absence of a clear ownership assignment.

The "anticommons" metaphor and, more generally, the existence of frictions in the market for intellectual assets, suggests that this last observation could be challenged, especially if we interpret the NPR approach as a theory of the allocation of intellectual ownership. As pointed out by Pagano and Rossi (2004), this should be considered the most appropriate interpretation of the NPR framework. This is mainly because the very meaning of specific investment is clear-cut when the assets at stake are subject to a form of legal protection (intellectual property law) that prevents their unauthorized reproduction. By contrast, the meaning of specificity is rather difficult to grasp when the assets considered do not enjoy this kind of protection and are therefore easily replicable.

According to Pagano and Rossi (2004), some of the NPR theory insights are thus applicable to the allocation of intellectual property rights as, for instance, the conclusion that ownership of an asset should be attributed to the agent with the most relevant investment in human capital specific to that asset and the prescription that complementary assets should be owned together. Interpreting the New Property Rights approach as a theory of intellectual ownership has, however, two less apparent consequences. First, it becomes unavoidable to relax the troublesome assumption made in GHM models that agents' coalitions are costlessy modified and optimal assets allocations are always achieved at zero transaction costs (thus allowing for the possibility of "anticommons" phenomena). Second, by taking into account the property of non-rivalness of intellectual assets, one is led to challenge from another perspective the NPR conclusion that that overall efficiency is enhanced by *any* property rights allocation, given that the very fact of making excludable a non-rival good implies an efficiency loss.

To understand the limits of the combination of markets and ownership of intellectual assets as a governance structure, one must thus take the New

Property Rights approach one step further and (a) introduce positive transaction costs in the exchange of assets; and (b) explicitly consider the dynamics of change of ownership allocation. The model we present in section 4 tries to accomplish these two tasks. The focus on dynamics represents one original contribution of this paper to the New Property Rights strand of research. The latter has been criticized on the grounds that it only considers equilibrium states (Franzini, 1993; Pagano, 1992). In our model, the dynamics implied by the two institutional arrangements we compare influence productive outcomes. However, before turning to the model, in the following section we introduce the idea that the "reverse anticommons" induced by the F/OSS licensing structure may be considered one form of private response to the poor performance of the combination of markets and ownership as a governance mechanism.

3 Free/Open Source Software as a "Reverse Anticommons"

F/OSS is defined as software that is (a) distributed not only in object-code form but also with the source code; and (b) distributed under a license that grants a number of rights to users, among which the right to freely modify it and the right to freely redistribute it. Thus, the true novelty of the F/OSS phenomenon is apparent from its very defining characteristic: the system of licenses under which it is distributed. The importance of the latter cannot be overstated. It is the particular system of licenses at the heart of F/OSS communities that sustains a development model characterized by quite peculiar characteristics, such as its distributed nature, the absence of an ex-ante centralized decision-making unit, the integration of users into production, parallel development and concurrent design and debugging.

Indeed, in contrast to what a first thought would suggest, F/OSS does not belong to the public domain (see, for instance, Gomulkiewitz, 1998; Perens, 1999). This is because developers do not surrender their rights to their software creations. All to the contrary, each contributor to an F/OSS project retains copyright over her work, while allowing free access to all those willing to accept the contractual provisions of a license compliant with the Open Source Definition (OSD)[7]. Were the software simply put into the public domain, everyone could appropriate it, modify and even obtain copyright over it and remove the author's name (Kennedy, 2001). By contrast, the combination of copyright and F/OSS licenses allows to strike a balance

[7]The Open Source Definition was published in 1998 as part of an effort by the Open Source Initiative to increase public awareness of the open source phenomenon. It defines the *"rights that a license must grant you to be certified as Open Source"* (Perens, 1999).

between the aim of ensuring the widest possible access to developers' intellectual creations and the possibility that the software be subtracted from public use. As McGowan (2001) puts it, *"[t]he licenses, and the GPL in particular, represent an elegant use of contractual terms and property rights to create social conditions in which software is produced on a model of openness rather than exclusion"*.

There exists a wide range of F/OSS licenses that differ primarily in regard to the way they address the issue of derivative works. A taxonomy of licenses could be build with reference to the latter criterion. At the more "permissive" end of the spectrum we would find licenses such as the Berkeley Software Distribution (BSD), the MIT and the Apache licenses, that do not impose any burden of reciprocity and allow the greatest freedom of modification, redistribution and license change. At the other end we would find the General Public License (GPL - also called copyleft license), which imposes that any derivative work that is distributed or published must be licensed as a whole under the terms of the same license (section 2(b) of the GPL license). This is often referred to as the "viral" character of GPL licenses, although it should be emphasized that it is not the mere use of GPL-ed software that gives rise to a reciprocal obligation to contribute back to the community one's own software code but only the development of works directly derived from or additional to GPL-ed software. In any case, the persistent nature of "copyleft" licenses must not represent a significant deterrent to their use if, as it is, they are the most widely used licenses in the F/OSS world (Lerner and Tirole, 2002).

In spite of the differences, all types of licenses share the trait of being the vehicle through which developers voluntarily suspend ex-ante, with respect to licensees, the exercise of their right to exclude from the intellectual assets they contribute to the community. This is one of the features of F/OSS licenses that, in our view, justify a definition of F/OSS as a "reverse anticommons". The property regime underlying F/OSS development is in principle akin to an "anticommons", as multiple individuals hold rights over complementary intellectual assets, thus giving rise to a web of blocking copyrights and therefore a situation of potential misalignment between use and exclusion rights. Against this background, the licensing structure at the heart of F/OSS development represents a way of reconciling use and exclusion rights because it creates a space in which exclusion rights are voluntarily suspended with respect to those that abide by the terms of the license. By accepting the terms of an F/OSS license users are granted a vast range of rights concerning the use, modification, redistribution etc. of the licensed software code while being prevented from excluding anyone from access to the code.

In other words, the F/OSS form of distribution constitutes an ex-ante contractual arrangement that turns a web of blocking copyrights into a situation in which - paraphrasing Heller and Eisenberg - *multiple owners have a right to use and no one has an effective privilege of exclusion.* Indeed, under F/OSS licenses, enforcement of the right to exclude from a given intellectual resource composed of many copyrighted contributions requires agreement among a sufficient number of right-holders to be effective. To put it differently, for a program licensed under a F/OSS license to be legally privatized in its current form, agreement among the multiple copyright holders should be reached. Transaction costs are thus incurred in the attempt to subtract the assets from public availability, rather than in the attempt to use them, as in a typical "anticommons" regime.[8]

Someone bored by the multiplication of definitions in the domain of property regimes might object that a "reverse anticommons" is nothing different than a "commons" property regime. The objection is understandable, as there are important similarities between the two property regimes: in both cases, ownership is not associated to exclusivity of use and owners can collectively exercise the right to exclude from a given resource. However, under a "reverse anticommons", such a resource is composed by a number of complementary assets to which each right-holder exercises exclusive rights[9], whereas under a "commons" a group of owners *jointly* enjoys exclusive rights to the resource. In addition to this, the contractual arrangements at the heart of the F/OSS "reverse anticommons" ensure complete freedom of use of the entire resource to all owners (and also to those non-owners that accept the terms of F/OSS licenses), whereas in a "commons" owners must collectively decide about the use of the resource. [10]

F/OSS can thus be looked at as a solution to anticommons problems, a solution that is perhaps less problematic than others that have been proposed, namely copyright collectives (Depoorter and Parisi, 2001), increased reliance on the fair use doctrine (Depoorter and Parisi, 2002), cross-licensing, patent pools and standard-setting organizations (Merges, 2000; Shapiro,

[8]Note that the fact that enforcement of the right to exclude from a given piece of software requires agreement among multiple right-holders and therefore implies significant transaction costs may generate problems also in circumstances in which agreement is necessary to protect the software from undesirable forms of appropriation. To avoid such problems, the Free Software Foundation encourages developers to transfer copyright to the foundation so as to facilitate enforcement (O'Mahoney, 2003).

[9]Of course, because of the adoption of F/OSS licenses, the right to use a given asset is not exercised exclusively, although in principle it could be.

[10]Another interesting difference between a "reverse anticommons" and a "commons" is that the first always arises as the outcome of a contractual arrangement, while the second may have such an origin (Eggertson, 1992; Lueck, 1994)but may also be the outcome of law, custom or regulation (Lueck and Miceli, 2004).

2002; Lin, 2002). What is interesting about this particular privately-devised solution to potential anticommons problems is that, differently from other solutions, it relies on a voluntary contractual reduction of the vector of rights that defines ownership.

The F/OSS phenomenon may thus be taken as a hint of the fact that in some circumstances the costs associated to ownership may outweigh its benefits. Such circumstances are particularly likely to arise when private ownership is defined over complementary assets and the degree of ownership fragmentation is high and, more generally, when the non-rival nature of the assets at stake implies the superiority of multiple access. Before formalizing these intuitions (section 5), however, we would like to add one last remark. The characterization of F/OSS as a "reverse anticommons" may suggest that Demsetz story about the emergence of property holds true but in a way its creator might not have intended: when the costs of internalizing externalities through property rights become too high relative to the benefits of internalization, a contraction of property rights may ensue. F/OSS is by no means the only phenomenon pointing in this direction. Similar instances of a voluntary renounce to exercise the right to exclude from a proprietary intellectual resource are observable, for instance, in the domain of so-called "open source biology" (Meurer, 2003) and in the context of high technology industrial districts (Gilson, 1999).[11]

4 A Stochastic Optimization Model of Property Rights

In this section we adapt a simple model from statistical physics to capture some of the issues that arise when the dynamics giving rise to a given allocation of property rights are examined. While detailed micro-modeling of agent choices and the internal decision making process of F/OSS projects and firms would be ideal, it would nevertheless obscure the issues we are looking at. This model is related to other random field models in the literature. Blume(1993) and Brock(1993) investigate similar statistical physics models applied to discrete choice games. Foley(1994) uses related concepts from thermodynamics to generate non-Walrasian general market equilibria. The basic idea of using a particular type of Markov process to optimize a complex function is known as simulated annealing, and a good exposition is

[11] Gilson has pointed out how firms in Silicon Valley voluntarily renounce to the possibility of opting out of a legal system that prevents covenants not to compete thus opening the way to significant knowledge spillovers and attributes to this characteristic of Silicon Valley the reason for the better performance of the latter over Route 128 industrial district. Thanks to Ron Daniels for calling our attention to this.

given in Azencott(1989).

As mentioned in the introduction, the model is also inspired by some insights derived from the New Property Rights approach and, in particular, by the view of the economy as composed of coalitions of agents introduced by Hart and Moore (1990). In the following two subsections (4.1 and 4.2) our focus will be on obtaining the traditional NPR solution and the "anticommons" property regimes as special cases of our dynamic model. The basic setting will thus parallel that introduced by Hart and Moore (1990): coalitions are to be interpreted as control structures, namely set of agents and the assets over which they have control. The reader can think about control structures as firms, as the NPR literature suggests although - again as in the NPR literature - we completely disregard transaction costs relative to the internal organization of the firm and therefore we are not, strictly speaking, providing insights directly relevant to the firm organizational form.

Let N be the set of agents (also the number of agents, in a typical abuse of notation). Agents' human capital endowments are heterogeneous and characterized by different degrees of complementarity with respect to other agents' human capital and to assets. Let A be the set (and the number) of assets. The function $V(x) = W(x_1, x_2, ...x_N) - \sum_{i=1}^{N} C(x_i)$ gives the total social surplus (see Hart and Moore(1990).

Agent i chooses x_i to maximize $v(x_i|S) - C(x_i)$ Where $v(x_i|S)$ is agent's i Shapely value given control structure S. This yields a function $x_i^*(S)$, of effort contributed as a function of the control structure. Agents' contributions will be lower than first-best, both because ownership can guarantee access to the asset, but not the cooperation of the other agents and because those excluded from ownership are exposed to the possibility of hold-up and thus have reduced incentives to invest. Still, the optimal ownership allocation allows to achieve second-best efficiency.

Define a state-space $G = 2^N \times 2^A$. A state in this space is a control structure: a set of agents and the assets over which they have control rights. We will define a stochastic process on this state space.

We impose the restriction that control structures can only change one member(agent or asset) at a time. That is, the only allowable transitions in each period are from X to $X \cup a$ or X/a where $a \in N \cup A$. Call this set of allowable destinations the neighbourhood of X, denoted N_X. This restriction represents a sufficiently reasonable approximation of reality and allows us to keep the model simple.

4.1 Costless Optimization

In this section, we offer an alternative formulation of a basic GHM insight: given zero transaction costs in the market for physical (in the present case, intellectual) capital, the allocation of ownership will achieve (second-best) efficiency.

To simplify the analysis, we assume that a transition from one control structure X to another Y happens with a probability that is a monotonic function of the payoffs: $P(Y|X) = \frac{\chi(N_X)}{|N|+|A|} min(1, e^{-(V(X)-V(Y))/C})$ if X $\neq$ Y.
This defines a Markov process. The parameter C is a measure of the weight attached to the payoffs in determining the transitions, and can be interpreted as the mutation probability: a low C means that agents are less likely to leave a relatively high-payoff state.

Thus a coalition can mutate its assets and agents with probability proportional to the gain in social efficiency. If there is a unique control-structure that optimizes V, the mass of the stationary distribution will be concentrated on that control-structure.

The ergodic distribution of this process has: $q_C(X) = \frac{e^{(V(X)/C}}{\sum_{Y \in G} e^{V(Y)/C}}$ (this is the probability of being at configuration X after a long time, also is the fraction of the time spent in state X). Note that $lim_{C \to 0} q_C(X) = 0$ unless $X \in argmax V$. Which means that the state that the system spends most of its time at is that which maximizes welfare, the traditional GHM solution.

Note that average payoffs are $\sum_{X \ inG} q(X)V(X)$.

4.2 Anticommons: Introducing Transactions Costs in the Exchange of Assets

This section introduces positive transaction costs in the above framework, so as to envisage the possibility of so-called "anticommons" tragedies.

Assume that $P(Y|X) = \frac{\chi(N_X)}{|N|+|A|} min(1, e^{-(V(X)-V(Y)-Z(X,Y))/C})$, where Z is a everywhere non-zero function capturing the costs associated with changing the control structure. Note that the probability of a transition from state X to state Y is inversely related to the overall number of agents and assets, which is consistent with standard results in the "anticommons" literature. If $Z(X,Y) = Z(X) + Z(Y)$(i.e. additively separable), then the stationary distribution, trivially enough, has $q_C(X) = \frac{e^{V(X)-Z(X)/C}}{\sum_{Y \in G} e^{V(Y)-Z(Y)/C}}$

Thus, there could be a Z that makes the system spend more time at a sub-optimal control structure. The average payoffs are again $\sum_{X\ inG} q_C(X)V(X)$. Note that the assumption of additive separability is quite strong: it implies that the transactions costs to get to X are independent of the current state Y.

4.3 F/OSS Optimization

This section addresses the core issue of the paper, namely the comparison between the F/OSS institutional environment and the more conventional governance structures based on the combination of ownership and markets referred to as "control structures" in the previous sections. Note that coalitions under the F/OSS institutional environment should be identified with "access structures", as the nature of F/OSS licenses implies that coalitions can only specify sets of agents and the assets to which they have access, rather than control. In order to pursue the comparison, we specify a number of assumptions.

Assumption 1: the payoffs to agents' investment are, *ceteris paribus*, lower under the F/OSS institutional environment than under the ownership-based institutional environment.

$$W^{F/OSS}(x) = \lambda W^{F}(x), \text{ where } \lambda < 1 \ .$$

This assumption reflects a first noteworthy difference between ownership-based governance structures and F/OSS communities. In the first case, the allocation of ownership (i.e. the right to exclude from access to an asset) can be used as an incentive mechanism but generates only second-best results because it exposes non-owners to the threat of hold-up. In F/OSS communities, agents' payoffs are, *ceteris paribus*, lower because of the classic public goods problem but, at the same time, agents face no reduction of incentives due to the threat of hold-up[12] . This is because participation to F/OSS projects is conditional on acceptance of licensing terms that prevent any of the agents from withdrawing their assets from the coalition.

Assumption 2: the mutation of coalitions in the asset space A involves greater transaction costs under the ownership-based institutional environment than under the F/OSS institutional environment.

[12]The extent to which agents' payoffs are reduced because of the public goods problem depends on the particular behavioral assumptions made. The reduction will be highest if one assumes rational self-interested agents and lower if other-regarding preferences are taken into account. As mentioned in the introduction, we do not deem it necessary to take a stance on this issue in the present paper, although we suspect that the issue of free riding in the context of F/OSS development is over-emphasized.

$$Z^{F/OSS}(X) > Z^{F}(X)$$

We refer here to the transaction costs relative to the exchange of assets across coalitions. Under the first institutional environment, exchange takes place through market-based mechanisms (purchase, licensing and mergers) that involve the kind of transactional hurdles pointed out in section 2. By contrast, under the F/OSS institutional environment, exchange of assets is governed by F/OSS licenses. One important aspect of such licenses is that they constitute a particular kind of mass-market or "shrinkwrap" licenses, characterized by the fact that the software user is presumed to be bound by the license as he downloads and uses the software, without the need for signing any license agreement (Kennedy, 2001). In other words, F/OSS licensing does not imply any bargaining and therefore transaction costs are reduced to a minimum. In addition to this, it should be noted that assets accrue to F/OSS coalitions also through voluntary contributions by the agents joining the coalition .

Assumption 3: the mutation rate C is consistently higher under the F/OSS institutional environment than under the ownership-based institutional environment.

$$C^{F/OSS} > C^{F}$$

The explanation for this hypothesis is twofold. First, in the conventional ownership and market-based institutional environment information concerning the nature of the human capital of agents inside and outside existing coalitions is conveyed mainly through prices (wages). The price system is not particularly good at conveying multidimensional information such as that concerning the extent of complementarity of the human capital of external and internal agents, potential increases in value due to the inclusion of an agent in a given coalition etc. The usual transaction costs associated mainly to problems of asymmetric information thus slow down the process of optimal coalition formation. By contrast, under the F/OSS institution the availability of the source code together with the public availability of other written information concerning each agent's contribution to the coalition allows to locate optimal coalitions much more quickly than under the price mechanism.

In addition to this, the "matching" literature (see for instance the survey by Mortensen and Pissarides, 1998) might suggest that, exactly because the modification of coalitions is associated to significant transaction costs and it is difficult for agents to find coalitions in which their human capital is most productive, once a "good match" is achieved, both the coalition and the agent are reluctant to break the relationship because the uncertainty

and the costs implied by the process of finding a new good match decrease both parties' outside options. By contrast, as there are less transaction costs involved in the recombination of F/OSS coalitions, developers will have less of a disincentive to leave a certain coalition and may more easily "jump" across different projects.

The expression for the probability of a transition from state X to state Y under the F/OSS institutional environment (with $\chi(N_X)$ the characteristic function of the neighbourhood of X) will thus be:

$$P(Y|X) = \frac{\chi(N_X)}{|N|+|A|} min(1, e^{-(\lambda(V(X)-V(Y))/C)}$$

Where $\lambda < 1$, reflects the fact that for a given access structure, F/OSS developers are unable to exploit the full joint-surplus produced, due to the fact that the software is non-excludable and non-rival. λ may be quite low. Note that the effect of a low λ is similar to the effect of a high C: it increases the probability of leaving a relatively low-payoff state.

Proposition: For an arbitrarily small λ, certain values of C and sufficiently high transactions costs in the market for the exchange of intellectual assets, F/OSS projects will outperform ownership-based governance structures, in the sense of higher average payoffs.

Proof: consider the ergodic distributions generated by the ownership-based institutional environment $q^f(Z)$ and by F/OSS, $q^{F/OSS}$, we can consider these vectors in $\mathbf{R}^{|G|}$.
We need to show that $< q^f(Z), V > \leq < q^{F/OSS}, V >$ for some Z.

If $V(X) - Z(X)$ has a different optimum, say X_{Z-opt} than V(X), then the result is straightforward, since we can pick C close enough to 0 s.t. $||q_C^{F/OSS} - 1_{X_{opt}}|| < \epsilon$ and $||q_C(Z)^F - 1_{X_{Z-opt}}|| < \epsilon$, where 1_A is the distribution that puts weight 1 on the control/access-structure A and 0 everywhere else. Since $V(X_{opt}) > V(X_{Z-opt})$ it follows that we can pick a C such that $< q_C(Z)^F, V > \leq < q^{F/OSS}, V >$.

If the transactions costs do not change the optimum, then the relative shape of V, λ, and Z matter. We can assert the following: if $\lambda V(X) \geq V(X) - Z(X)$ for all X, then F/OSS has a higher mean payoff for a sufficiently low C.

The result stated in the above proposition can also be expressed differently: in spite of the incentive effect ownership generates (captured in the model by the parameter λ), it might be the case that a different institutional environment that does not rely on such incentives can achieve *on average*

better productive outcomes, in circumstances in which transaction-cost induced frictions in the market for intellectual assets tend to give rise to an "anticommons" situation. The intuition behind this result is the following: there is a trade-off between the stability of coalitions and the probability that they reach a global optimum. The terms of this trade-off differ under the two institutional environments we compare. In particular, access structures (under the F/OSS institutional environment) have an advantage with respect to control structures (under the ownership-based institutional environment) in searching the space of possible configurations, although they tend to be less stable than control structures.

5 The Legal Protection of Software and the F/OSS development model

In this section, we would like to advance the hypothesis that our model could provide some intuitions as regard to the question of the opportunity of an increased reliance on patents as the prevailing form of legal protection of software. In this regard, harmonization of legislation at the international level has not been reached (yet). In the United States, the presumption of unpatentability of software-related inventions was gradually eroded by a series of court decisions, beginning with the 1981 Supreme Court decision that held that software that was part of a manufacturing product or process was to be considered patentable (*Diamond vs. Diehr*). Patent protection became thus available for software creations along with the two traditionally adopted means of protection: trade secrecy and copyright. In the European Union the debate opened by the proposed "directive on the patentability of computer-implemented inventions" is still unsettled. As of today, no agreement has been reached on the opportunity of a convergence towards US standards, although the activity of pressure groups opposing software patentability seems to have been effective so far.

The debate on software patentability parallels, to some extent, the more general confrontation between "transactional pessimists" and "transactional optimists" that was mentioned in section 2. "Transactional pessimists" tend to emphasize that the strengthening of intellectual property protection for software can increase the likelihood of the occurrence of a "tragedy of the anticommons" in the software domain by deepening the wedge between use and exclusion rights. "Transactional optimists" tend, on the contrary, to stress the fact that software patentability might contribute to the creation of a market for software components, as argued by some commentators (Lemley and O'Brian, 1997), because a clearer definition of the boundaries of IP rights facilitates transactions (see also Merges, 1999).

As for the first hypothesis, numerous features of both software patents and the nature of software innovation represent a cause of concern. First of all, it should be noted that the degree of property rights fragmentation is likely to be particularly high in software, both because of the fine granularity of the patents granted (Kahin, 2003) and because of the heterogeneous nature of the players in the software industry. Further, valuation problems are likely to be substantial in the case of software patents, as the difficulty of ascertaining the existence of prior art in the absence of comprehensive databases and expert examiners generates a high degree of uncertainty on validity. Combining these two insights with the observation that software products are complex artifacts that result from the assembly of a high number of components in the context of an innovative process characterized by a high degree of cumulativeness and complementarity, one is led to take seriously the possibility that an extreme fragmentation of property rights could stifle software innovation. This possibility appears even more concrete when one considers that complex technologies constitute the domain most prone to the strategic use of patents, as shown by Cohen at al.(2002) in the context of complex technologies different from software.

Bessen and Hunt (2003) offer some evidence pointing in this direction, in an often cited study that explores patenting behavior in software and the relationship between patenting and R&D spending in the US. Their main results reveal that the substantial drop in the cost of obtaining patent protection has resulted, beginning from the 1990s, in a dramatic increase in software patent applications and that patents tend to substitute, rather than complement, R&D activity. The latter result is particularly interesting, because it supports the hypothesis that firms increasingly recur to "patent thicket" strategies. In other words, it supports the idea that firms build large patent portfolios in order, on one side, to bargain to obtain access to externally developed technologies and, on the other side, to defend themselves from infringement suits through the threat of counter-suits. It should be noted, however, that Bessen and Hunt's paper has been repeatedly criticized, especially from the methodological point of view (Hahn and Wallsten, 2003).

The second hypothesis is supported by the belief that copyright protection of software has traditionally encouraged the inefficient practice of reinvention and has hampered the development of a market for tradeable software components. Indeed, insofar as copyright only protects the expression of a given idea, rather than the idea itself, it allows developers to appropriate others' innovative creations as long as they refrain from literal copying of the software code. The trend toward increased reliance on patent protection - the argument goes - may reduce the extent of reinvention, in-

crease software reuse and favor the development of a market for software components. This is because patents provide a strong form of protection and because patent law *"precludes the overlapping claims to interface elements that copyright law produces"* (Lemley and O'Brian, 1997, p).

We do not have conclusive evidence on the likelihood of either of the outcomes sketched out above, although we tend to be sympathetic with the "pessimist" stance. What we want to emphasize here is that, in evaluating the merits of patent protection for software, its effects on the F/OSS phenomenon should be taken into due account.

If there is some truth in the intuition that the diffusion of the F/OSS development model can mitigate to some extent "anticommons" problems, it becomes particularly important to assess the impact of the recent trend toward increased reliance on patent protection for software on F/OSS. One possible cause of concern in this regard is that the diffusion of software patents may erode F/OSS coalitions' ability to acquire assets. Assets contributed to F/OSS communities may increasingly become the object of infringement litigation, as the number of patents granted raises with few limits and firms rely progressively more on patents as strategic weapons. What is more, F/OSS communities lack both financial resources and bargaining chips in the form of accumulated patent portfolios to successfully resist infringement actions. What this means is that the diffusion of software patents may put at risk the survival of a form of innovation that has demonstrated so far a high degree of vitality.

We thus join the chorus of those commentators that have already warned about the negative incidence of software patents on F/OSS (Benkler, 2002; Bessen, 2002; Lessig, 2002), suggesting that this aspect should be appropriately taken into account in the overall assessment of the desirability of this form of IP protection for software creations. Our model suggests that, if the dismal forecast of "transactional pessimists" is correct, patent protection of software might be pushing toward "the worst of both worlds": a situation in which the efficiency of ownership and market-based governance structures is reduced by the rise in transaction costs due to patent-induced extreme property rights fragmentation and, at the same time, the productive possibilities of F/OSS communities are constrained by a contraction of the accessible asset space.

6 Conclusion

In this paper, we address an issue that has never been the object of formal treatment: how does the performance of governance mechanisms based on the combination of ownership and markets and F/OSS communities compare? We focus on the dynamics of change over time of the control/access structures under the two institutional environments and we point out the existence of a trade-off between the stability of a coalition under a given institutional environment and the probability that it reaches a global optimum. The model shows that, in spite of the incentives ownership provides, under certain values of the mutation rate and sufficiently high transaction costs, F/OSS projects outperform ownership-based mechanisms, i.e. they generate higher average payoffs.

We argue that one of the strenghts of the F/OSS institution is that the system of licenses on which it is based creates a "reverse anticommons", a situation in which transaction costs for the exchange of assets are reduced to a minimum. This constitutes a crucial advantage especially in circumstances in which governance structures based on the combination of ownership and markets are characterized by a tendency toward the occurrence of "anticommons" tragedies and, more generally, when there are significant transaction costs in the market for intellectual capital, as it seems to be the case in the software domain.

The model leads us to question the wisdom of increased reliance on patent protection for software. Such a shift in the legal framework for software protection may end up obtaining "the worst of both worlds": it may lower the performance of ownership and market-based mechanisms by increasing the extent of property rights fragmentation and therefore transaction costs and it may jeopardize the development of F/OSS communities by reducing the number of assets F/OSS developers may lawfully have access to.