To Patent or Not to Patent: A Pilot Experiment on Incentives to Copyright in a Sequential Innovation Setting

Paolo Crosetto

To Patent or not to Patent: a Pilot Experiment on Incentives to Copyright in a Sequential Innovation Setting

Paolo Crosetto[1,2]

[1] DEAS, Università degli Studi di Milano, paolo.crosetto@unimi.it
[2] Dipartimento di Scienze Economiche e Aziendali, Libera Universita Internazionale degli Studi Sociali (LUISS), Roma, paolo.crosetto@gmail.com

Abstract. This paper presents preliminary results from a pilot experiment dealing with the economic motivations to contribute to Free/Open Source Software (FOSS). Bessen and Maskin [1] argue that in a dynamic sequential innovation framework the standard argument for granting patent protection is no more valid and the innovator has at certain conditions an incentive to fully disclose the results of his works; in these same conditions, a copyright strategy could result in a *tragedy of the anticommons* [5, 2].
We study in the lab the choice of copyrighting or copylefting subsequent innovations in a dynamic setting *à la* Bessen and Maskin, introducing an innovative experimental design requiring real effort on the part of subjects. The players are asked to actually 'innovate' producing words from given letters, and face the choice to *copyright* or *copyleft* their words. Preliminary results show that copyleft is more likely to emerge when royalty fees are relatively high, and when the extendability, modularity and manipulability of inventions is enhanced.

1 Introduction

The standard economic argument in favour of patents and copyright states that, since knowledge is a public good with a high fixed cost of production and a relatively small cost of imitation, the state needs to grant for limited time a monopoly on the invention to allow the inventor to recover his costs, thus giving him an incentive to invest in the first place. While a long-standing debate in economics exists over the exact nature of the cost-benefit tradeoff at the base of intellectual property regulations from society's point of view, it has been widely assumed that an innovator would copyright/patent his innovation whenever given the chance to do so. This assumption seems to need further inquiry as we witness the rise of FOSS.

It has recently been suggested in the economics of innovation literature that the very nature of software - encoded, algorithmic knowledge - could generate incentives to disclose the creation of one's labour and ingenuity. It has been shown by [1] that when innovations are sequential and research efforts are complementary, in a dynamic setting it is optimal for innovators to fully disclose their work,

since every inventor gains from the marginal future contributions of others to his own innovation, made possible by the open nature of his contribution.

The converse argument, i.e. that in dynamic, sequential settings introducing intellectual property can slow down innovation, has been made several times in economics [10], and has recently been the focus of the literature about the so-called *tragedy of the anticommons*. The name derives from the famous tragedy of the commons paper by [4]: while in the commons the absence of exclusive usage rights generated overutilisation and waste, in the anticommons the presence of overlapping and fragmented exclusion rights generates underuse and waste. The imposition of patents on research results generates an anticommons whenever every right holder independently sets a price for the license without taking into account negative externalities, with the result of general stagnation of downstream innovation. This mechanism has been outlined theoretically [2, 9], analysed experimentally [3] and proved to be at work in the field of biomedical research [6].

In this paper we present the preliminary results of a new real-effort experimental design that enables to explore jointly, in a controlled setting, the intellectual property choices of subjects in a dynamic sequential innovation game and the aggregate consequences of individual behaviour in terms of anticommons.

2 The model in Bessen and Maskin [1]

When innovations are sequential, each innovation is both an output and an input for further innovations developed by different inventors; the true value of the first invention might only be revealed when the original idea is extended, in directions possibly unforeseen by the first innovator.

In this context, copyright has ambiguous effects, from both the society and individual points of view, as Intellectual Property is likely to affect not only the revenue, but also the costs of the innovator. Copyright generates an expected stream of revenues for the copyright holder, thus providing incentives to put effort into the (costly) innovating activity. At the same time, follow-on innovators face the cost of licensing and increasing transaction costs if bundling of many innovations patented by different inventors is needed, thus generating a negative effect on the number of follow-on innovations.

A possible - if radical - solution to the anticommons problems in R&D would be for the innovator to release its discovery under copyeft licences ('private-collective' innovation mode, [11]). Nonetheless, while it is straightforward to see why a 'private-collective' solution could spread once initiated - innovators can freely use and improve upon increasing amounts of knowledge - it is harder to see why a self-interested innovator should voluntarily forgo any direct appropriation of revenue from his discovery, releasing it free to the public.

According to [1], the very sequential nature of innovations can provide incentives to start sharing and to forgo direct appropriation through IP. In their theoretical paper patents are shown to be a valid policy tool - enhancing social welfare as well as providing incentives to innovate to all firms involved, thus

increasing the likelihood of discovery - in a static model; when the same model is enlarged to include dynamics and sequential and complementary innovations, though, patents have a counterintuitive effect, and are shown, in line with the anticommons literature, to become *hurdles* to innovation rather than *spurs*, the social welfare being higher *without* patents than with them. More to our point, Bessen and Maskin show that if profit dissipation due to increased competition is low - i.e. if the availability of further innovations expands the market, increasing opportunities for further R&D on the one side and increasing the number of interested customers on the other - the firms themselves have an incentive *not to* patent their discoveries even if a patent system is available at no cost to them.

The present paper presents preliminary results of a new experimental design developed to test in the lab the argument of [1].

3 Experimental design

The experiment has been designed to create a dynamic, interdependent setting in which subjects make choices in a collectively and dinamically co-created landscape. To simulate innovative activity, the experiment follows a real effort protocol; the players have to actually innovate - over a set of given rules - using both economic (experimental money) and cognitive (creative effort) resources.

The real effort task chosen is a sort of *Scrabble* game. The idea of using creation of words to mimic innovative activity in investigating FOSS has been pioneered by Lang et al. [7], who design a double-auction market for words, inducing demand but leaving players free to make supply decisions. The design here presented differs greatly from theirs, their interest being in market design, while the main focus here is on sequential innovations.

The design chosen creates a situation in all respects similar to the [1] model, but it is not a formal equivalent of the latter; rather, it is a transposition of the assumptions and workings of the model into an intuitive but controlled setting.

3.1 The *copyleft* game

1. *Game structure.* The game is played by **6** players. Players play in turns until there are no letters left in the letter set, with random assignment of starting positions. The players aim to maximise their monetary payoff from the game. Players get payoffs from two sources: 'use value' of the words they produced or extended, and '(net) royalty fee revenue'. The game interface is shown in Figure 3.
2. *Letter set.* The letter set is composed of **200** letters (see Figure 1). Frequencies of letters are the same as in standard *Scrabble*. Every letter comes with an attached payoff this letter will give the player when inserted into a produced or extended word.
3. *Turns.* In any turn, a player can make at most three actions: buy a random letter (for a fixed price); either produce a word or extend an existing word; decide whether to copyright or copyleft the produced word. Players

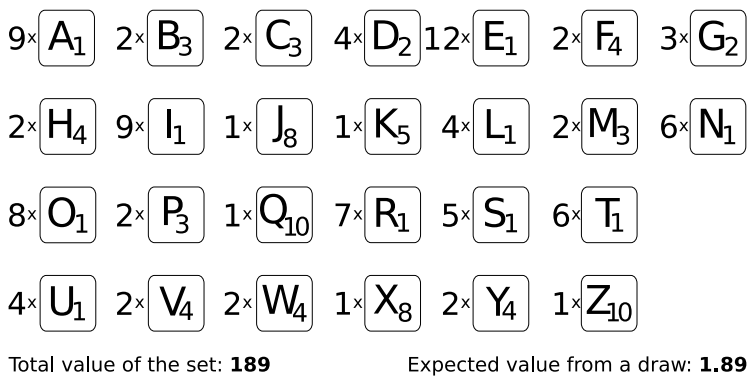

Fig. 1. Letter set used for the game

can choose to make just one or two decisions, or to pass. Turn structure is summarised in Figure 2.

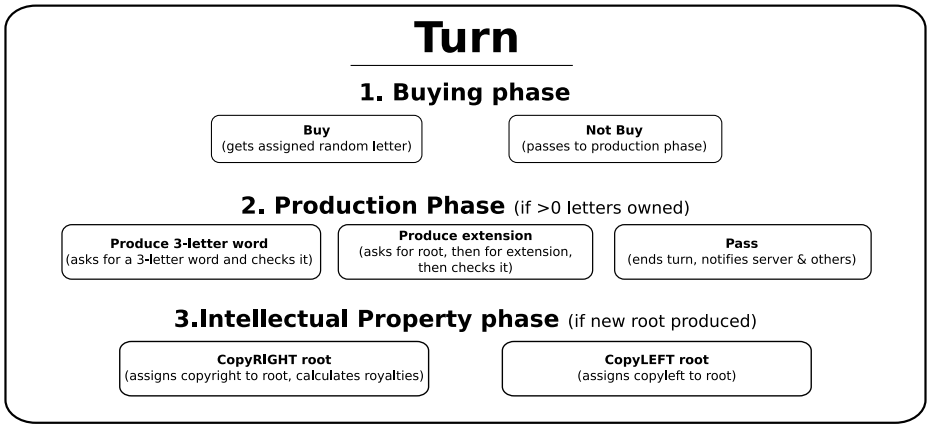

Fig. 2. Actions available to the players in every turn

4. *Buy phase.* When it is their turn, players can buy a random letter for a fixed price, out of their show-up fee. If a player buys, a letter is drawn from the letter set (without substitution) and assigned to the player. The price is set at 2 experimental points, which is slightly above the expected value from the set, 1.89.
5. *Words phase.* Players collect payoffs by composing or extending words. A spelling check is performed on produced words. At every turn, every player can perform only one action in the words phase, either producte a new or extend an existing word.

- *Producing words.* Players may produce from scratch only three-letter words, called *roots*. This rule mimics the fact that an innovation must be more-or-less working when released. Production of a word generates a 'use value' payoff that is the sum of the values of the letters used to form the word. This payoff is incurred only once, the moment the root is produced. Produced words are common knowledge; the use of the root as input for extensions is instead subject to the IP choice of the creator.
- *Extending words.* Players may extend existing words, by one letter at a time, inserting this letter in any position within the word or at its edges. Anagrams are not permitted. Words have no length constraints apart from the ones implicit in the English language. Example extension paths for the root *car* are *car* → *car**e*** → ***s****care* → *scare**d***, or *car* → *car**d*** → *card**s***. Extending a word gives as payoff the sum of all the letters composing the extended word but might be subject to copyright fees to be paid to the word original creator.

6. *Spellchecker.* To avoid language skills biases, the players are provided with an interactive free-access spellchecker embedded in the main game interface.
7. *Word availability.* Extended words do not replace existing roots, but exist in parallel to them. A root can be used by different extenders to generate different word trails.
8. *Intellectual Property Phase.* Players that have produced new roots must decide whether to copyright or copyleft the word or extension. The IP choice cannot be undone in later stages and lasts throughout the game.
   (a) *Copyrighting.* Copyright can be obtained free of charge, and gives the copyright holder the exclusive right over the use of his root/extension. This implies that no other player can produce the same root/extension and that the copyright holder will receive a royalty fee every time some player uses it for an extension. The fee is fixed, and is proportional to the value contributed to the word. A root which value is v generates a fee of αv, in which $0 < \alpha < 1$ is a fixed known parameter. Every extension falls under the same rule, and if the value added to the word is w, an extension generates a right to obtain a royalty fee of αw. The n^{th} extender of a word of value $v + w$ pays $\alpha(v + w)$ in fees, distributed automatically to other players according to the value contributed. Free copyright (an irrealistic assumption) has been introduced to strengthen the external validity of the emergence of copyleft in the lab.
   (b) *Copylefting.* Copyleft is free of charge. It gives the copyleft owner no exclusive rights, but it endows users with a large set of use, redistribution and modification rights. Extenders can use the word as an input for free, but they are not allowed to copyright their extension. When extending a word, a copyleft extender earns all of the use value of the word, without having to pay any fee, but not getting any further fee from extenders either.

The game described above has been designed to share as many features as possible with the model in [1], but it is not a perfect match. Every deviation from

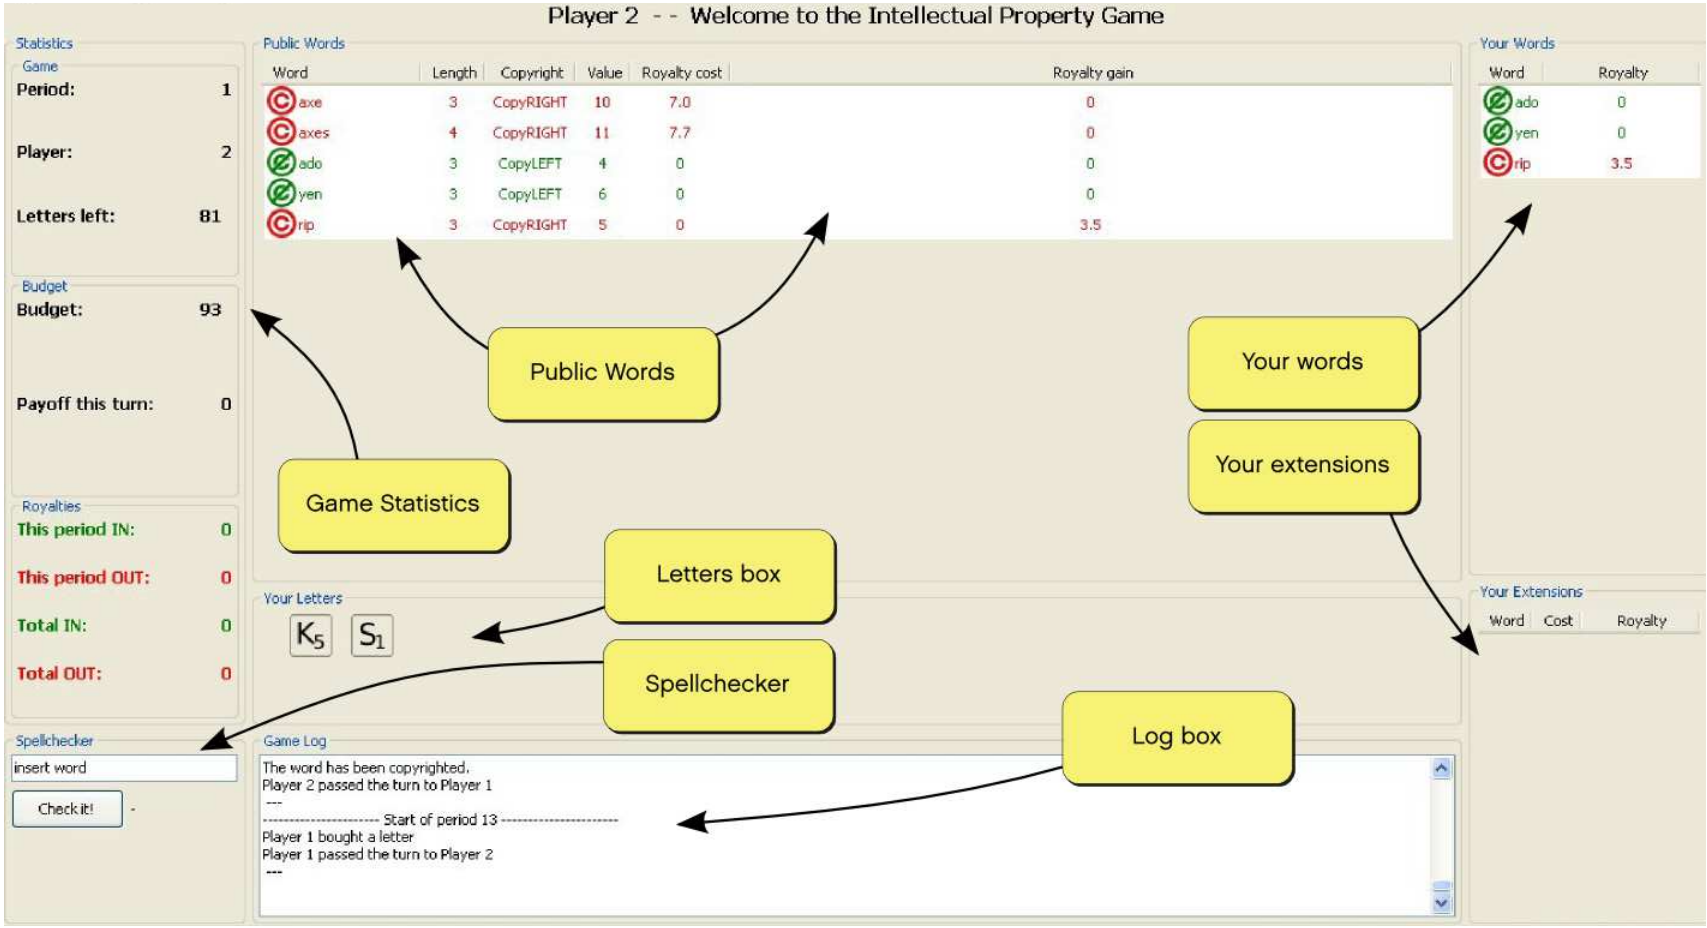

Fig. 3. The main interface for the Copyleft Game

the model was made on the safe side, lowering the likelihood that the Bessen and Maskin outcome would emerge.

4 Experimental results

The pilot experiment was held in October 2009 in the EXEC lab, University of York, UK. The software was developed in Python [8]. The experiment involved in total 23 English native speakers. Four sessions of 6 players were run, divided into two treatments: in the *low-fee* treatment the royalty parameter was set to $\alpha = 0.7$, while in the *high-fee* treatment it was set to $\alpha = 0.9$. We expected the *low-fee* session to show less copyleft activity and anticommons problems than the *high-fee* sessions.

Copyleft emerged in three out of four sessions (the two *high-fee* and one of the *low-fee*), dominating one; anticommons features (underuse of resources, low cooperation, lower payoffs) emerged in the two *high-fee* sessions as expected.

4.1 General results

The sessions differed from one another in many respects. Two main problems for data analysis emerged: first, Session 4 featured only 5 players instead of the standard 6; second, in Session 3 the players found a clever way of altering the game rules to their advantage, and hence played a completely different game. Safe comparisons can be made between Session 1 (*low-fee*) and Session 2 (*high-fee*). Results from Session 3 will be separately discussed (Section 4.4).

The experiment showed significant emergence of copyleft in the *high-fee* treatments, and a number of recurring features.

1. The number of roots created is 30 in all sessions. What differentiates the sessions is the number and nature of extensions; the behaviour of follow-on innovators determines social welfare and payoffs.
2. Payoffs are significantly higher in the *low-fee* Session 1 w.r.t. the *high-fee* Session 2 (Mann-Whitney U, p-value = 0.041, K-S p-value = 0.1429) and less so w.r.t. Session 4 (Mann-Whitney U, p-value = 0.129, K-S p-value = 0.5909).
3. The *high-fee* treatment sessions resulted in a lower number of words created, and in a higher waste of productive resources (unused letters), resulting in lower average payoffs. In *high-fee* sessions players tended to extend their own roots and to avoid other player's roots, to access which they had to pay royalties; this behaviour created quite isolated word trails and lower extension opportunities.
4. Copyrighted roots are much more likely to be created than copylefted roots. Players did not create any (but one on the very last period) copylefted root in Session 1 (*low-fee*), but did create 5 and 1 copylefted roots in Sessions 2 and 4 (*high-fee*)

	Low-fee		*High-fee*	
	Session1 (6)	**Session3 (6)**	**Session2 (6)**	**Session4 (5)**
Payoff: average (£)	17.76	26	13.95	15.7
Standard Deviation	2.6	4.8	1.95	2.52
Total number of words	112	154	84	91
% copyrighted	99.1%	29.2%	73.8%	90.1%
Total net value created	471	968	237	288
Royalty turnout	268.8	71.4	119.7	71.1
Royalties/value	32.2%	5.3%	21.6%	11.6%
Unused letters	26	16	46	47
Number of copyrighted roots	30	11	30	30
Extensions per root	2.7	3.1	1.06	1.73
Relative extensions per root	0.19	0.27	0.1	0.12
Number of copylefted roots	1	4	5	1
Extensions per root	0	26.25	3.4	8
Relative extensions per root	0	0.4	0.42	0.34
Average word length	4.18	4.53	3.76	3.89
Copyrighted words	4.18	3.73	3.58	3.8
Copylefted words	3	4.86	4.27	4.67
Average word value	7.45	8.63	6.6	6.76
copyrighted words	7.46	6.84	6.22	6.62
copylefted words	6	9.37	7.68	8.11

Table 1. Summary data for the four experimental sessions

5. Copylefted roots are three to seven times more likely to be extended than copyrighted roots. Once copylefted words exist, the incentives to extend them are higher, because the foregone royalty from one's extension is more than counterbalanced by the fact that no royalties are paid on the root.
6. Copylefted roots generate a two to four times higher exploration of extension possibilities (see 'Relative Extensions' rows in Table 1). The 'average relative extensions per root' index was computed as the average of the ratio of the actual number of extension of each root and the theoretically possible number of extensions of each root allowed by the English language.
7. Copylefted words are on average longer and have higher value than copyrighted words. This is because copylefted words are extended more and new roots tend to be copyrighted but not extended: when faced with the choice of copyrighting or copylefting a root, the player most of the time decides to copyright it; but this decreases the likelihood of follow-on innovations to be built on top of that root, and hence slows down the pace (extensions per root), and reduces the exhaustivity (relative extensions per root) of further inventions.
8. The amount of royalties paid or received is higher in *low-fee* sessions. When fees are low, players tend to seize the best opportunities available to them irrespectively of the fact that there is a royalty to be paid: this results in a high number of words, lower letter waste, and a high amount of royalties exchanged. When royalties are higher (*high-fee* treatment) players tend to restrict their innovative efforts to the 'free' roots, either copylefted or owned by the extender, and hence many opportunities are missed, a higher number of letters is left unused and a lower amount of royalties flows in the system.
9. Players used the spellchecker extensively - on average more than 100 times - in all sessions. There is no significant correlation between the use of the spellchecker and the result in terms of payoff for the player.

A treatment effect can be argued to exist. Copyleft emerged and saw sustained contributions in *high-fee* sessions, while it did not emerge in *low-fee* session; anticommons features appeared in (*high-fee*) sessions, in which payoffs were lower, number of unused letters higher, and royalty flows reduced.

The paradoxical effect that better patent protection results in a lower amount of innovations, one of the central points in [1] is reproduced in the lab; moreover, players recognise the mild incentives to copyleft, and provide some (though not many) copylefted roots; once a copylefted root exists, players reap much more benefits from it, extending it further and deeper, than from copyrighted roots. There is not enough evidence, though, to support the stronger argument in Bessen and Maskin's paper that in a sequential settings the firms prefer not to patent their innovations, relying on subsequent innovations rather than on royalties from upstream contributions: if that were the case, we would have recorded a surge in copylefted *roots* as the game progressed. What happened instead was that in all sessions players preferred to release the *roots* created as copyright, even when this generated a suboptimal anticommons situation; nonetheless, players preferred to extend existing copylefted roots much more

than they extended copyrighted roots, and hence copylefted words appeared and accounted for up to a quarter of all words created.

4.2 Session 1 (*low-fee*): copyright and *business as usual*

Six subjects participated in Session 1, a *low-fee* treatment. In the session 112 words, all copyrighted but one, were produced: on average, every player produced 18.7 words, 3.6 words per period. The only copyleftd word was created in last turn by the last player, when the choice was of no importance any more.

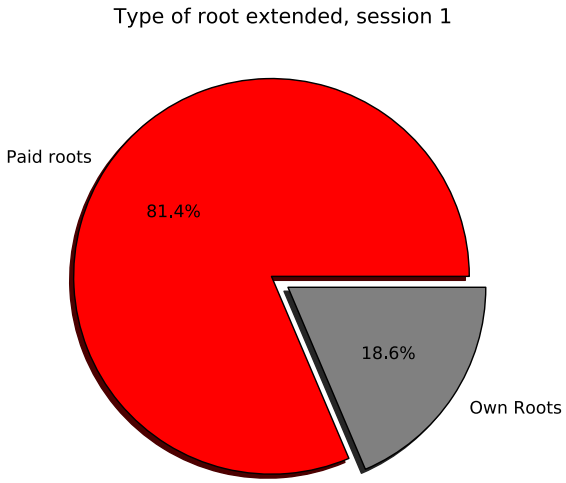

Fig. 4. Distribution of extensions by root type, session 1

The session describes a 'business as usual' situation. Players copyrighted every root created and pursued the most profitable extensions without taking into account the royalty fee. 81.4% of the extensions came from extending a root created by other players, and hence incurring the royalty fee; conversely, only 18.6% of extensions were built upon own words (Figure 4). Since a player is likely to produce one sixth of the roots, this figure implies that, on average, the players extended any root available and had no biases favouring their own creations.

Royalties make up an important share of players' payoffs. In the session the net added value created was 471 experimental points (worth 47.1£); royalties exchanged were worth 268.8 points, i.e. 57% of the value added.

The effect of royalty flows on individual payoffs can be appreciated in Figure 5: some players produced a lot of added value, but paid high amounts in royalties to do so, while other players produced fewer - but more likely to be extended - roots and extensions, and received a consistent amount of royalties. The two strategies are somewhat complementary: some inventors prefer to create many, if marginal, innovations; others focus on seminal, if fewer, innovations.

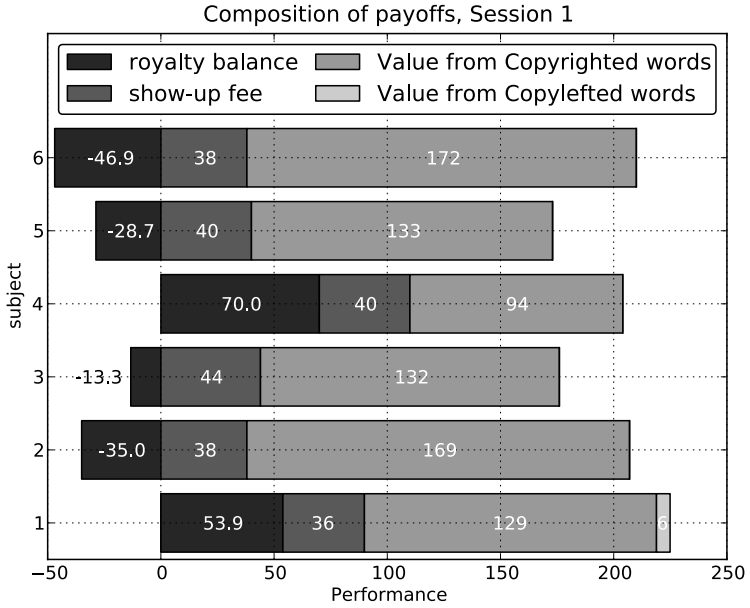

Fig. 5. Breakdown of payoffs, session 1

The session showed an intellectual property system working well: copyright does not get in the way and royalty fees are considered low enough not to block follow-on innovations. Copyright generates considerable incentives to innovate rewarding the creators of the best innovations (in this case, the most extendable) through royalties and the hard-working 'marginal' inventors through use-value of their marginal inventions. No 'tragedy of the anticommons' appears, as downstream innovators are not blocked in their endeavour by excessively high royalties, upstream inventors can trust that their roots will be used, and royalties flow from the marginal inventors, extending existing technologies, to the basic inventors, creating promising roots. In an ideal situation of perfect information, with no transaction nor legal costs associated with IP and with low, fixed and ex-ante known royalty fees, the copyright system works and delivers incentives to different types of innovators.

4.3 Session 2 (*high-fee*): copyleft and anticommons

Six players took part in Session 2, a *high-fee* treatment. The players left 23 letters unbought and 23 more unused. The players created 84 words, of which 22 copylefted: 14 words per player and 2.5 per period. The session showed both the emergence of copyleft and of anticommons gridlock: players in Session 2 created less words and enjoyed lower final payoffs w.r.t. Session 1 (see Table 1); moreover, the letter waste was double (46 *vs* 23 unused letters), the net value

added half (237 *vs* 471), the average length (3.76) and value (6.6) of the words produced was substantially lower than in Session 1 (4.18 and 7.45).

73.8% of the words created were copyrighted, 26.2% copylefted. The players created 32 copyrighted extensions from 30 roots (1.07 per root), and 17 copylefted extensions from 5 roots (3.4 per root, three times as much). Copylefted words are furthermore longer and more valuable (see Figure 6). The best efforts of the players were directed to extend copylefted words. The players showed a somewhat schizophrenic behaviour: when creating new roots, players opted for copyright, hoping to enjoy a royalty stream from extenders, but, when extending, preferred to extend either their own or copylefted roots (Figure 8, right). This resulted in a mild but significant 'tragedy of the anticommons'; the players could have broke out of it either by adopting a 'business as usual' strategy as seen in Session 1 (but the high royalty fee seems to have made this less likely) or by starting to create copylefted roots. Having failed to do both, the players ended up producing less words and earning lower payoffs.

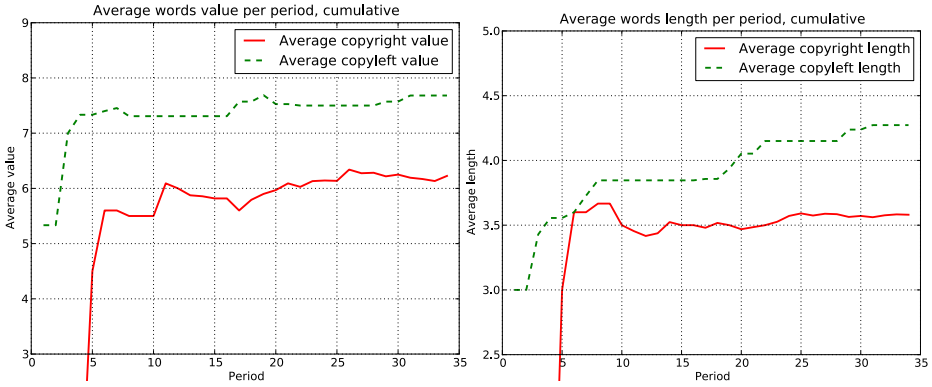

Fig. 6. Average word value and length per period per type, session 2

Players started copylefting roots at the very beginning, and switched to copyright after the first 'defections', replicating the usual 'decay of contributions' phenomenon of Public Good games (Figure 8, left). After period 8 all newly created roots were copyrighted; nonetheless, copylefted roots continued to be extended during all the rest of the game, possibly to avoid paying royalty fees to the other players.

The players tended to extend other player's roots much less than in the *low-fee* treatment: only 51% of the extensions came from a root owned by another player, down from 81.4% in Session 1 (Figure 8, right). This resulted in many opportunities for extension being lost.

In Session 2 anticommons effects appeared, breaking the flow of innovations and resulting in lower value added and lower payoffs. Awarding a higher royalty to innovators generates the perverse effect of lowering the amount of innovation

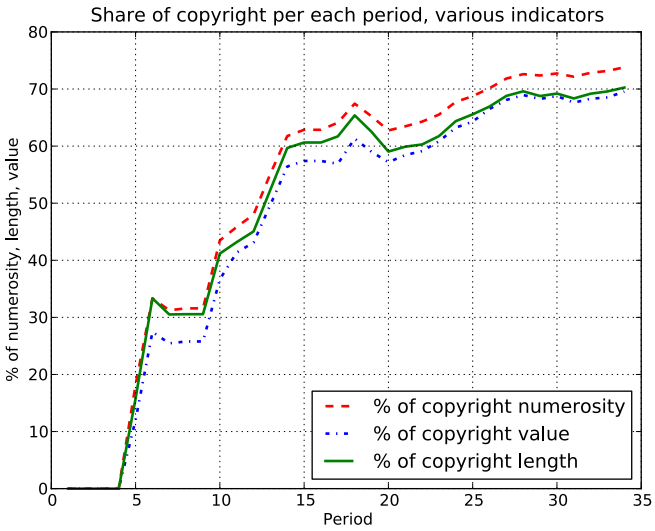

Fig. 7. Share of copyrighted words, length and value per period, session 2

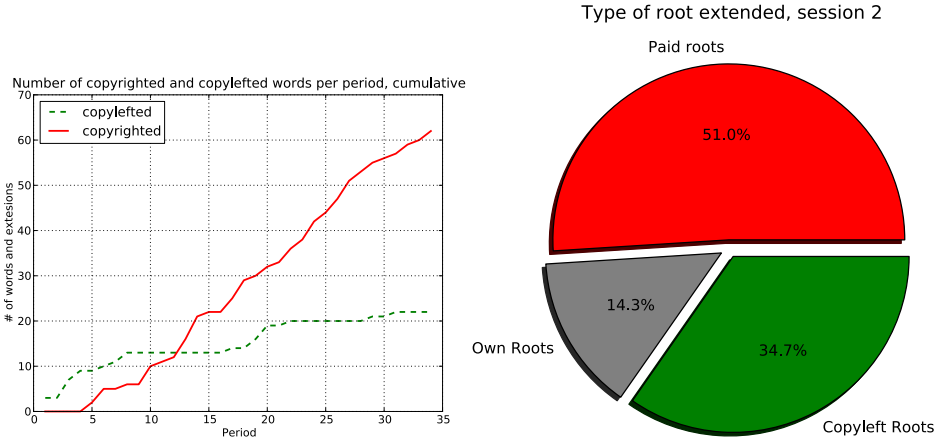

Fig. 8. Session 2: Words created by period and by type (left) and extensions by root type (right)

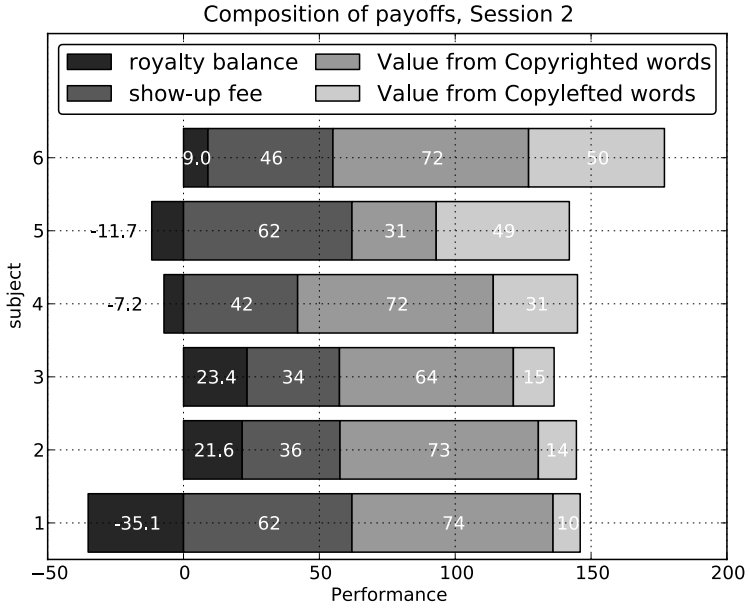

Fig. 9. Breakdown of payoffs, session 2

created in the system, at the same time not benefiting the inventors. Upstream patents, coupled with high royalty fee, discourage follow-on inventors to innovate on top of someone else's basic research on the one side, and to invest *tout-court* on the other hand (unused letters). Contrary to expectations, though, players fail to see the incentive to get rid of the patent system altogether, as argued by [1]: they extend copylefted roots but do not forgo the (low, given the behaviour of players) expected stream of royalties from copyright. This fact could lead to argue that while incentives to copyleft exist, they imply a much more forward looking attitude than the one displayed by subjects in this session.

4.4 Session 3 (*low-fee*): *learning to cheat* and copyleft

Six subjects participated in Session 3, a *low-fee* treatment. The session is radically different from the others, because of cheating behaviour on the part of the subjects that was not foreseen and hence not stopped by the software. When the cheating behaviour was noticed, it was decided to let the session proceed to see what this 'innovative' behaviour would mean in terms of aggregate statistics.

Early in the game, a player accidentally found a bug in the software: she extended '*new*' into '*wine*', and was not stopped by the system. The software was not designed to check the order of the remaining letters of the root, thus leaving the door ajar for a particular kind of anagram, generated by adding one letter in any position and displacing the other letters at will. Actually created examples of this anagrams are given by *lady* → *delay*, *preen* → *opener*. All other

types of extensions were correctly turned down by the software, including simple displacement (with no added letter, e.g. *not* → *ton*) and inclusion of more than one letter (e.g. *sit* → ***s**h**i**f**t***).

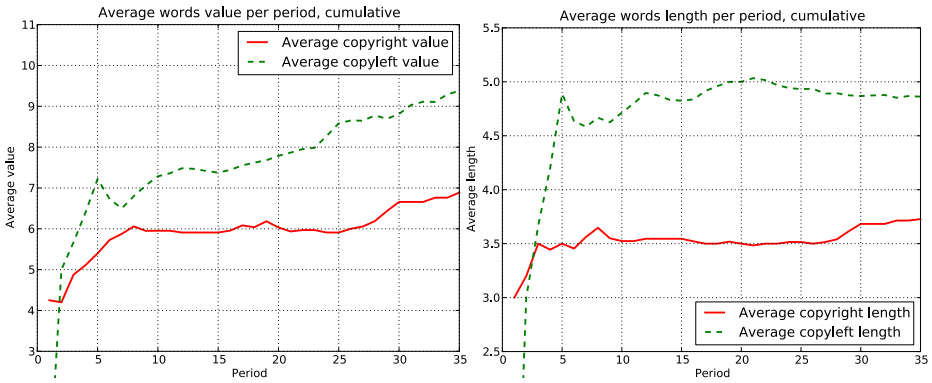

Fig. 10. Average word value and length per period per type, session 3

This new rule implies higher tinkering possibilites and higher payoffs, but at much higher cognitive costs, and it allows for a highly enhanced extendability of words, thus giving less incentives to create new roots.

After the first player successfully anagrammed a root to form an extension, the other players noticed and, all but one, endorsed the new practice. Around period 20, five out of six players were exclusively using the new rule. A further consequence of this endogenous discovery of value-enhancing rules was that the players eventually abandoned copyright completely: given the much enhanced extendability, players found it much more profitable to copyleft the word and then build on the contributions of the other players than to rely on royalty fees that were unlikely to be received anyway.

The results of the new rules were striking. In the session 154 words were created, 25.6 per player and 4.5 per period; the net value added reached the level of 968, twice the level of Session 1, and four times the level of Session2; the royalty flow accounted for a mere 7.4% of it, and the unused letters were just 16, the lowest value of the experiment. The average value and length of copylefted words increased during the experiment much more than those of copyrighted words (see Figure 10), and the percentage of copyrighted words dropped steadily, ending at 29.2% (see Figure 11).

In Figure 12, left, it can be seen the point, around period 10, when the new rule began to spread. The creation of new copyrighted roots did not stop altogether, but continued at a much lower pace. The breakdown of payoffs records the importance of copyleft for the very high final payoff enjoyed by subjects (Figure 13). The number of extensions generated using a copyrighted root (see Figure 12, right) dropped to 12%, as did the use of own roots; three-quarters of

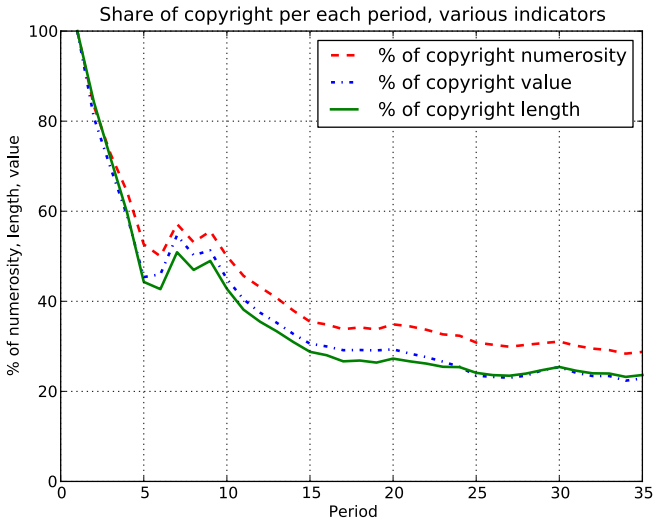

Fig. 11. Share of copyrighted words, length and value per period, session 3

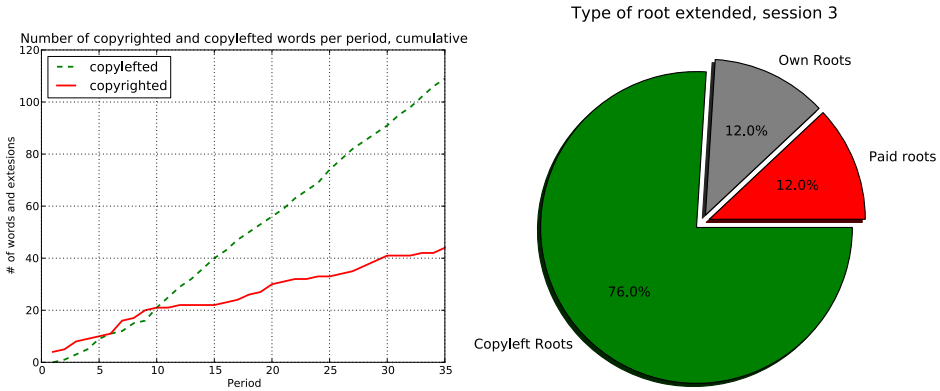

Fig. 12. Session 3: Words created by period and by type (left) and extensions by root type (right)

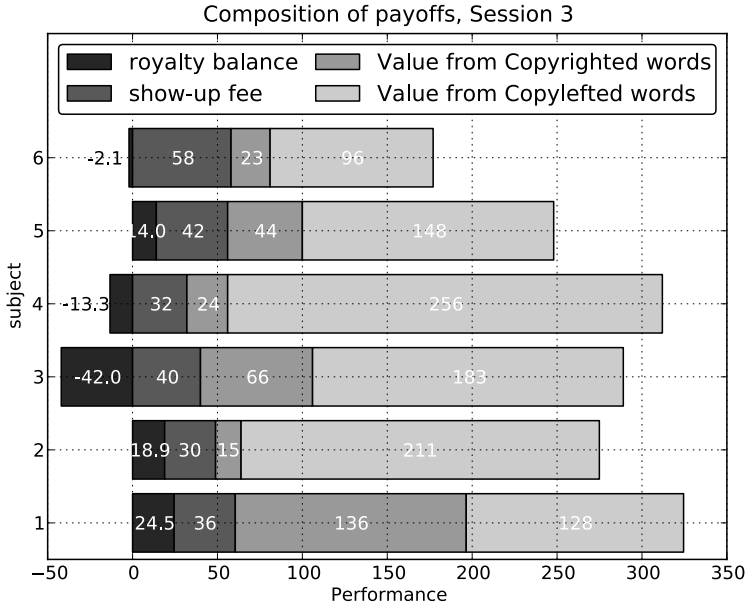

Fig. 13. Breakdown of payoffs, session 3

extensions were built on top of the existing 4 copylefted roots, with an impressive 26.25 extension for each copylefted root.

The session showed two important facts: first, the players learned, from observation and trial and error, a superior set of rules and adopted it quickly; second, when extendability is enhanced the incentives to copyleft are much higher, and the role of copyright is more limited.

4.5 Session 4 (*high-fee*): copyleft and anticommons

Only five subjects participated in Session 4, a *high-fee* treatment. The players created 91 words, of which 9 copylefted; 18 words per player and 2.67 per period. The session had many results in common with Session 2, the other *high-fee* session: the players were affected by the high royalty fee and preferred to extend their own roots rather than pay royalties to the other players (Figure 14, right); this resulted in a substantial letter waste (47 unused letters), low payoffs (average 15.7£), a low number of words produced (even though higher than in Session 2), low average word length (3.89) and value (6.76).

In this session copylefted roots were extended deeper and further than copyrighted ones, but the players failed to find a way out of the anticommons: all roots but one were copyrighted, despite the fact that players were not likely to extend other player's roots (just 28% of extensions came from a paid root). The 'isolation' of players was the highest of all sessions, with 58% of extensions created from one's own set of roots (Figure 14, right).

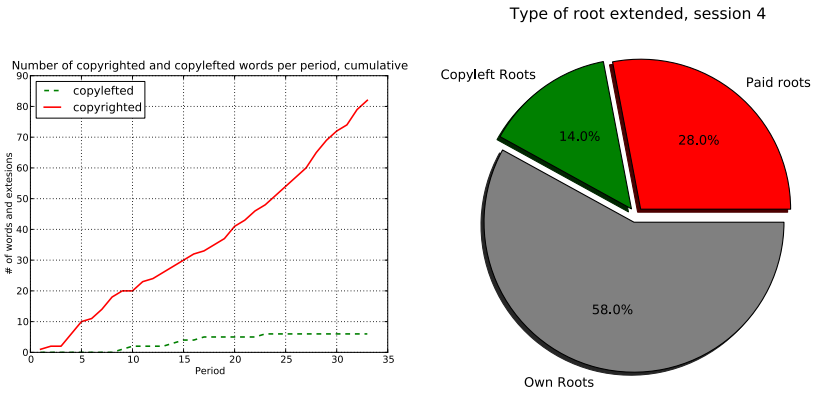

Fig. 14. Session 4: Words created by period and by type (left) and extensions by root type (right)

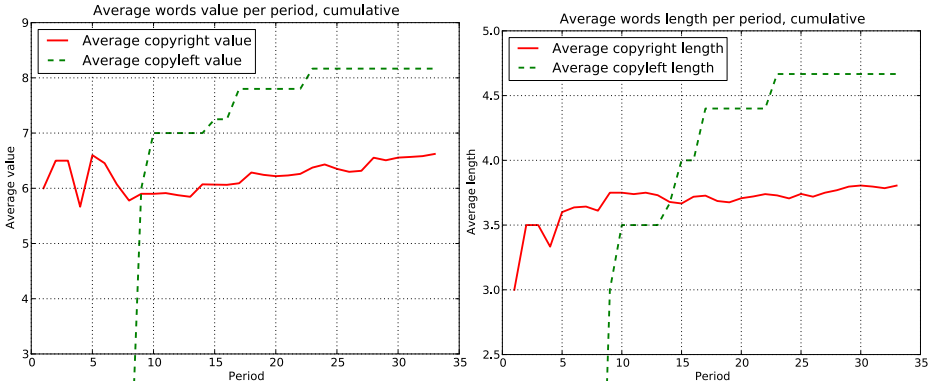

Fig. 15. Average word value and length per period per type, session 4

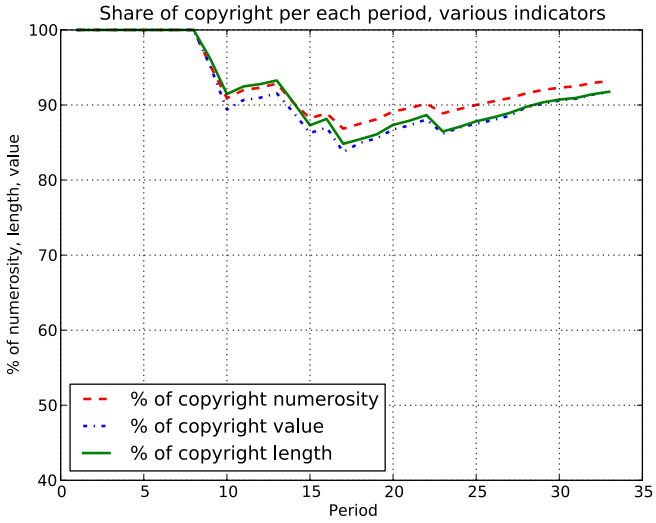

Fig. 16. Share of copyrighted words, length and value per period, session 4

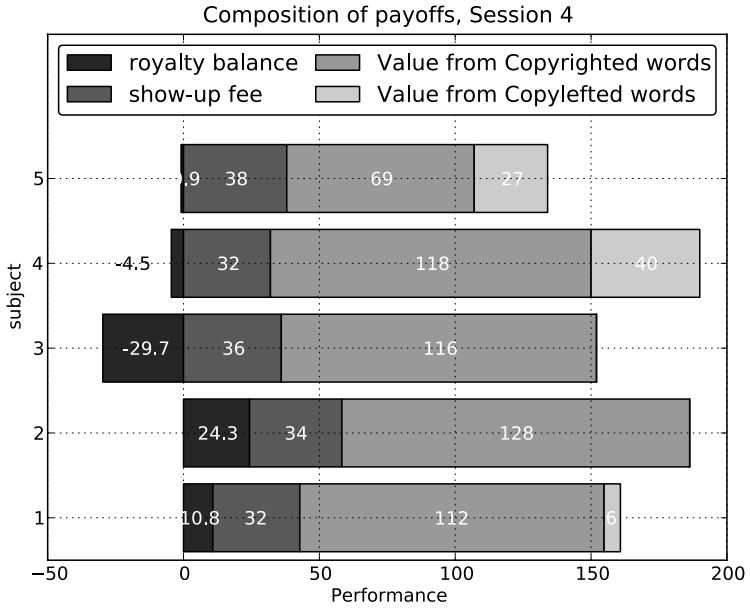

Fig. 17. Breakdown of payoffs, session 4

As in all other sessions, the copylefted words showed a higher average value and length than the copyrighted ones (see Figure 15); here they represented just 10% of the overall value, though, and had a minor impact on payoffs (Figure 17). Royalties also played a minor role: the total net value added was 288, with royalty flows at a low 71.1, 24.7% of the added value. This was due to the high reliance of players on own roots rather than on cross-fertilisation of word trails with other players.

Session 4 confirms that when royalty fees are high, players fail to build innovation on top of other players' roots, reducing the number of profitable opportunities. Copyleft is a possible way out of this gridlock situation, but players fail to recognise this fact.

5 Conclusions

This paper sets up a dynamic real-effort experimental game to test the argument of [1] that in a dynamic setting featuring sequential and complementary innovations the innovators themselves would, at certain conditions, prefer not to impose copyright on their creation and instead welcome imitation.

Even if the statistical validity of such a small experiment is admittedly very low, preliminary results show that both anticommons problems and copyleft are more likely to emerge when royalty fees are relatively high. Increasing intellectual property rights leads to less innovation and to potential gridlock on the one hand, and to firms giving up copyright protection and embracing 'open' innovations. Copyleft emerged and saw sustained contributions in *high-fee* sessions, while it did not emerge in the *low-fee* session; anticommons features appeared in *high-fee* sessions, resulting in less innovations created, lower payoffs, higher number of wasted resources (on the form of unused letters), and reduced royalty flows, w.r.t. *low-fee* sessions.

The paradoxical effect that better patent protection (higher fees) results in a lower amount of innovations is reproduced in the lab. There is not enough evidence, though, to support the argument that in a sequential settings the firms prefer not to patent their innovations, relying on subsequent innovations rather than on royalties from upstream contributions.

The experiment is also potentially able to account for the gradual engagement in the production of Open Source software of more and more corporate giants. In the experiment, while it is hard to find someone willing to start a copyleft trail, the incentives to contribute to the copyleft commons with extensions once such an alternative exists are high. This leads, in time, to the building of a sizeable body of copyleft roots and extesions. As this grows, more and more subjects - even the ones that would strongly prefer copyright - have incentives to extend the copyleft words, because they take out more and more value from the commons. The copyleft rule forces them to give away any rights on their marginal contributions; but, as the size of the copyleft commons increases, the cost-benefit balance of contributing to it shifts gradually to the benefit side. The copyleft clauses are able to create a growing commons; no matter why the

first contributors decided to forgo their copyright fees, the existence of a free alternative is an attractor for a larger and larger number of developers.

The preliminary experimental evidence calls for further exploration of this experimental design, both deepening - i.e. with more replications of the same treatments, to gain statistical significance - and widening the research agneda. The word-game design could be used as a platform for testing the effect of different intellectual property policies or different technological settings on the amount of innovative activity, licensing choices, social welfare, commons and anticommons problems. The experimental design would allow for instance to test the consequences on innovation levels and quality of varying the length and breadth of patent protection, or varying the menu of licenses available, or allowing patent pools.

References

1. Bessen, J., Maskin, E.: Sequential innovation, patents, and imitation. Economics Working Papers 0025, Institute for Advanced Study, School of Social Science (Mar 2006), http://ideas.repec.org/p/ads/wpaper/0025.html
2. Buchanan, J.M., Yoon, Y.J.: Symmetric tragedies: Commons and anticommons. Journal of Law & Economics 43(1), 1–13 (April 2000), http://ideas.repec.org/a/ucp/jlawec/v43y2000i1p1-13.html
3. Depoorter, B., Vanneste, S.: Putting humpty dumpty back together: Experimental evidence of anticommons tragedies. SSRN eLibrary (2004)
4. Hardin, G.: The tragedy of the commons. Science 162(3859), 1243–1248 (1968), http://www.sciencemag.org/cgi/content/abstract/162/3859/1243
5. Heller, M.A.: The tragedy of the anticommons: Property in the transition from marx to markets. Harvard Law Review 111, 621–88 (1998)
6. Heller, M.A., Eisenberg, R.S.: Can patents deter innovation? the anticommons in biomedical research. Science 280(5364), 698–701 (1998), http://www.sciencemag.org/cgi/content/abstract/280/5364/698
7. Lang, K.R., Shang, R.D., Vragov, R.: Designing markets for co-production of digital culture goods. Decision Support Systems 48, 33–45 (2009)
8. van Rossum, G.: Python reference manual. CWI Report (May 1995)
9. Schulz, N., Parisi, F., Depoorter, B.: Fragmentation in property: Towards a general model. Journal of Institutional and Theoretical Economics 158, 594–61 (2002)
10. Scotchmer, S.: Standing on the shoulders of giants: Cumulative research and the patent law. Journal of Economic Perspectives 5(1), 29–41 (Winter 1991), http://ideas.repec.org/a/aea/jecper/v5y1991i1p29-41.html
11. Von Hippel, E., Von Krogh, G.: Open source software and the 'private-collective' innovation model: Issues for organization science. Organization Science 14, 209–223 (2003)