Journal of Economic Behavior and Organization 206 (2023) 204–221

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Intellectual property reform in the laboratory

Ismaël Benslimane[a], Paolo Crosetto[b,*], Raul Magni-Berton[c], Simon Varaine[d]

[a] Univ. Grenoble Alpes, IPhiG, Grenoble 38000, France
[b] Univ. Grenoble Alpes, INRAE, CNRS, Grenoble INP, GAEL, Grenoble 38000, France
[c] ESPOL-LAB, Université Catholique de Lille, Lille, F-59000, France
[d] Univ. Grenoble Alpes, Sciences Po Grenoble, PACTE, Grenoble 38000, France

ARTICLE INFO

Article history:
Received 29 March 2022
Revised 5 December 2022
Accepted 7 December 2022

JEL classification:
034
D9
D72

Keywords:
Intellectual property
Patents
Institutional reform
Vote
Laboratory experiment
Innovation policy
Real effort task
Creativity

ABSTRACT

This study experimentally captures the effects of reforming intellectual property (IP) and measures whether abolishing IP after a vote can reduce the creativity of the most talented innovators. The subjects start in a baseline setting with IP and then move to a second unannounced phase in which IP has or has not been abolished. First, we manipulate exogenously the presence or absence of the possibility to vote for or against IP. Second, we manipulate the information given to the subjects after the first phase. In the information treatment, we inform the subjects right before the vote that a previous experiment that used the same design (Brüggemann et al., 2016) showed that the no IP regime significantly increased players' gains. Contrary to preregistered expectations, the results show that undergoing a vote does not reduce overall creativity. Actually, the most talented innovators do not vote in favor of IP. Rather, the subjects who vote in favor of IP are those who benefit relatively more from royalties. Surprisingly, no correlation is found between these two populations: the IP in our experiment seems to not reward the best players, but the players who choose an 'IP-driven' strategy of focusing on more extensible words while simultaneously relying on their own creations forego cross-fertilization with other players. These are relatively low-skilled players who choose a rent-seeking strategy that maximizes gains *from* the IP system *itself*. There are plausible arguments that this result is at least partly valid in the real world, especially for complex and highly sequential innovations where it has been proven that patent trolls and anti-competitive strategies are important.

1. Introduction

Institutional change is difficult – even when the change is known to be beneficial. Changing the institutions that govern the functioning of society is likely to generate short-term disruption, increase transaction costs, and cause errors and inefficiencies. Moreover, there might be losers – people who do not benefit from the change, either at all or in the short run. These losers could reduce their contributions to society because of their grievances against the winners and their feeling of relative deprivation to thus limit or erase the benefits of the institutional change (Dannenberg and Gallier, 2019).

Reforms of intellectual property (IP) rights, particularly patents, are a case in point. Despite a broad debate, the net effect of intellectual property on innovation and welfare is unclear. It is thus even more unclear whether a reform would

* Corresponding author.
E-mail address: paolo.crosetto@inrae.fr (P. Crosetto).

https://doi.org/10.1016/j.jebo.2022.12.005
0167-2681/

be beneficial. Already back in 1955, the U.S. Senate Judiciary Committee commissioned several reports from experts such as Vannevar Bush and Fritz Machlup on the effects of the patent system and the need to reform it. After surveying 200 years of economic theory on the patent system, Machlup (1958) concluded the following:

> *No economist, on the basis of present knowledge, could possibly state with certainty that the patent system, as it now operates, confers a net benefit or a net loss upon society. [...] If one does not know whether a system 'as a whole' [...] is good or bad, the safest "policy conclusion" is to "muddle through" – either with it, if one has long lived with it, or without it, if one has lived without it. If we did not have a patent system, it would be irresponsible, on the basis of our present knowledge of its economic consequences, to recommend instituting one. But since we have had a patent system for a long time, it would be irresponsible, on the basis of our present knowledge, to recommend abolishing it.*

According to Machlup, in the 1950s, we did not know enough about the advantages and disadvantages of the patent system to take risks involved in its demise. Despite decades of research, Machlup's conclusion still stands – we do not know enough about the effects of the patent system to be able to suggest a clear reform path. In the meantime, the stakes involved in taking action have steadily increased. In Machlup's time, less than 150,000 patents were filed worldwide in a given year. This number increased 20-fold to over 3 million in 2018.

Given the current size and impact of the IP system, are IP abolitionists such as Boldrin and Levine (2008a) or IP reformists such as Bessen and Meurer (2008) and Stiglitz (2008) playing with fire by suggesting that we should abolish or strongly reform IP? IP is now long established and has attained global reach. The disruption risk is therefore greater. Economic agents accustomed to the IP system could be destabilized, demotivated and, in the worst case, not compensated by the path opened by IP reforms.

As Ouellette (2015) pointed out, many scholars have addressed the first-order question in patent law: What policies should we adopt to promote innovation? However, given the vast empirical uncertainties about patents, more attention should be given to the second-order question: What policies should we adopt to promote innovation about promoting innovation? Ouellette (2015) calls for "*patent experimentalism*": that is, to experiment with reforms in the field, target certain technologies, or create legislative variety across different jurisdictions. The second-order question must also take into account the obstructions that most reforms produce.

In this paper, we contribute to the second-order question in an exercise of *patent experimentalism*. We conduct a laboratory experiment to test Machlup's fear of eliminating a system such as IP, which has existed for a long time and whose effects are still not clearly understood. We exploit the possibility of building counterfactual situations in the laboratory to investigate the effects of *endogenous institutional change* on innovative activity. We observe subjects' behavior in a sequential innovation setting as they transition from a default system with IP to a system without IP.

We adopt the experimental design of Crosetto (2010); Brüggemann et al. (2016) and create a sequential innovation setting in the laboratory by using a *Scrabble*-like word-creation task. The subjects invest in letters and can use them to create words that yield them a direct payoff as in *Scrabble*. Moreover, the subjects can produce new words by extending other subjects' words after the payment of a royalty to the original innovator. This design allows us to easily and intuitively build a sequential innovation setting in the laboratory and to study institutional change by abolishing the royalty due to the extension of other subjects' words. We choose this setting because it has been shown to yield *better* outcomes *without* IP. This feature allows us to focus on how the institutional setting and details of the *transition* impact behavior and welfare even if we know that the change is for the better.

Within this setting, we focus on the political nature of the transition and its potential effects.

First, we let the subjects vote. In the `Vote` condition, the subjects play once in a context with IP and then face a vote to decide whether to abolish or keep IP in a second, unannounced round. The default policy is the abolition of IP; this can be overturned by a unanimous vote to keep IP in place. This default+veto voting rule allows us to maximize the observations in which we can study institutional change and, crucially, to create many losers whose behavior interests us the most.

Second, we provide the subjects with information before voting about the results of the Brüggemann et al. (2016) paper. In the paper, when playing a similar *Scrabble*-like game, groups with no IP consistently and substantially outperformed groups with IP. The idea of conveying scientific information about the relative performance of the two systems mimics an information campaign that could accompany the institutional change.

As a control, we run treatments where the subjects are not allowed to vote and just play once in the context of IP and once, without announcement, with no IP. Since there is no vote, the subjects have no way to act on any information that we might give them. Nevertheless, to generate the cleanest possible counterfactual, in the off-chance that the information had a *direct* effect on play in the second period, this control is split between groups that receive information about the results of the Brüggemann et al. (2016) paper and groups that do not receive this information.

In a publicly preregistered analysis plan (details on OSF: https://osf.io/nzq4f), we hypothesized that the democratic nature of the transition would paradoxically *lower* the beneficial impact of abolishing IP.

Proving that one system is better than another does not tell us much about the effect of the *transition* between systems. The hallmark of public policy change is that subjects can compare different systems and as a result, experience relative deprivation in adopting one system over another. Even if the after-reform system is shown to yield a positive impact on social welfare in vitro, a coalition of losers from the institutional change could try to block the change and after losing the battle, adopt boycotting behavior that would make the switch not cost-effective. There is ample evidence, e.g., in ultimatum games (see Güth and Kocher, 2014, for a review), that people pay an individual price to send a message, enforce a moral or social norm, or show their anger.

In the context of the IP in democratic societies that we study here, an institutional change toward the abolition of IP could occur following a democratic procedure (e.g., a general election), and innovators could feel victims of a "tyranny of the majority" that wants to expropriate them.

Assuming that the main losers of an IP ban are the more talented innovators, because they are those who benefit more from IP, we hypothesized that banning IP produces grievances among the most productive group and leads them to reduce their effort in sequential innovation. We thus expected the beneficial effect of the abolition of no IP to be higher when players do not vote and to be lower in the presence of a vote.

We further hypothesized that providing information about the likely outcome of the policy change might ease the transition. Even losers would be less inclined to boycott the change if they know from a trusted external source that the transition will bring in net benefits. We accordingly expected the problem identified above to be less harsh in the presence of information.

Testing these hypotheses has consequences on the procedure by which IP should be amended. When majority preferences oppose the preferences of the best innovators, the social welfare could be reduced. This is one of the arguments to enshrine IP in the constitution so that the "procedures for change require a lengthy ratification process or something more than simple majority approval" (Anderson and Hill, 1988).

The results show that our hypotheses were partly off the mark.

We replicate – with a lower significance – the positive effect on innovation when abolishing IP found in Brüggemann et al. (2016). Removing IP is beneficial to all players, irrespective of their previous productivity and their status as losers or winners in the IP system. The few groups that unanimously voted to keep IP performed worse than the groups that abandoned IP. This means that the chosen setting delivered what we needed – i.e., that abolishing IP is beneficial in this experiment.

However, contrary to our expectations, introducing a vote *increases* the positive effect of abolishing IP. This is due to an unexpected finding: the subjects voting to keep the IP system are not the most talented innovators but those who have earned more from royalties. Surprisingly, no correlation is found between these two populations: in our experiment, the patent system seems not to reward the best players, but the players who choose an "autarkic" strategy of relying on their own creations and forego cross-fertilization with other players. These are relatively low-skilled players who choose a rent-seeking strategy that maximizes gains *from* the IP system *itself*.

Providing subjects with the right to vote is not a threat for the best players because their preferences are not homogeneous, and if there is a trend in their vote, it is to lean slightly *against* IP. As a consequence, groups that unanimously vote to keep the IP are more likely to be composed of rent-seekers without particular skills, and they fare relatively worse. Providing information has a large effect on the likelihood of voting to abolish IP but has no effect on innovation and welfare, further showing that the mechanism that we hypothesized – good players are squeezed out of their profits and retaliate by reducing their effort – did not materialize.

Notably, the hypothesis that failed is so widely believed to be trivially true as to border on common sense. One of, if not *the*, main argument of the proposers of IP rights is that they provide incentives to innovators and that without such incentives, we would see less innovations, and everyone would be worse off. However, at least in our experimental setting, IP does not benefit the most productive innovators but rather the agents who are best at gaming the incentive scheme that it provides. Given the large evidence on patent trolling, the large and growing profits of legal firms that provide patent services, and the widespread practice by key innovators to choose trade secrets rather than use the IP system, there are plausible arguments that our results are at least partly valid in the real world.

2. Related literature

Whether patents and intellectual property in general increase public welfare is a debated and not obvious question. Patents solve a dynamic inefficiency – by anticipating being copied and exploited in the future, innovators do not devote effort – by legally creating a static inefficiency – by granting a temporary monopoly on the invention. Their impact on R&D, innovation and overall welfare is therefore theoretically unclear and depends on the exact assumptions of each model (for the case of sequential innovation, see Scotchmer, 1991; Green and Scotchmer, 1995; Moschini and Yerokhin, 2008).

Considerable research has generally studied monetary incentives on creativity (Kaufman and Sternberg, 2019). However, these studies provide mixed results, where some find positive effects (Eisenberger and Rhoades, 2001; Eisenberger and Shanock, 2003), others find positive effects only on certain aspects such as quantity and quality but not on originality (Charness and Grieco, 2014; Laske and Schroeder, 2017), others find no particular effect (Eckartz et al., 2012) and still others find negative effects of incentives (Ariely et al., 2009; Erat and Gneezy, 2016).

There is a growing literature in experimental innovation economics (for a review, see Buccafusco and Sprigman, 2019). Although the experiments vary widely in scope and context, most papers develop tasks that mimic innovation as a search over some large, multidimensional space unknown to the subjects but controlled by the experimenter. Most of these studies (Buchanan and Wilson, 2014; Buccafusco and Sprigman, 2010; Erat and Gneezy, 2016; Eckartz et al., 2012) do not find any effect of IP on innovation or creativity, and some even find negative effects compared to other incentive mechanisms (Meloso et al., 2009; Charness and Grieco, 2014). Brüggemann et al. (2016) find that IP rights in a scrabble-like sequential game significantly decrease innovation because the introduction of IP leads to more rudimentary innovations, as subjects do not exploit the most promising sequential innovation pathways.

There is extensive literature on biases in favor of the status quo. Making choices is often difficult, and decision makers may prefer the easier option of maintaining the status quo. In this case, the cognitive costs of decision making may outweigh the benefits of a superior choice. As evidence, decision makers are more likely to postpone a decision as alternatives are added (Tversky and Shafir, 1992), and the preference for the status quo increases with the number of options (Samuelson and Zeckhauser, 1988).

The literature distinguishes the decision maker's preference not to act (omission bias) from the preference for maintaining the status quo, but both appear to be based on loss aversion and regret avoidance (Eidelman and Crandall, 2012). Thus, individuals give more weight to losses than to equivalent gains, and there is also greater regret for action than for inaction. Since the status quo serves as a reference point from which change is considered, the costs of change outweigh the potential benefits, and since regret is stronger for action than for inaction, the two mechanisms combined create a relative advantage for the current state of affairs (Moshinsky and Bar-Hillel, 2010).

Eidelman et al. (2009) have also shown that people consider the mere existence of a thing as proof of its goodness and have a preference bias for things that have been around for a long time (Eidelman et al., 2010).

Our work is the first to study the impact of IP *democratic reform* on innovation outcomes in the laboratory. Buchanan and Wilson (2014) experimentally tested the effect of an exogeneous change from a no IP treatment to an IP treatment on innovation and found no effect on innovators' investment. Our study assesses the effect of an endogenous IP reform through a vote – that mimics the way that the IP regime would change in democratic countries and raises specific questions. Additionally, close to our study, Sprigman et al. (2013) designed an experiment to determine whether subjects' valuation of attribution was stable against changes in the default rule that governed the availability of attribution of a creative work (in this case, a painting).

3. Experimental design

We implement a *Scrabble*-like word creation game in the lab. The design replicates most features of the one employed in Brüggemann et al. (2016). The game is played in groups of three players who are randomly matched. They are assigned letters and then take turns in creating simple (three-, four-, five-letter-long) words – that generate a payoff as in Scrabble, which is roughly determined by the inverse of a letter's absolute number in the set – or extend existing words by adding one and only one letter in any position. Extensions generate a similar Scrabble-like payoff, and all the letters of the words are counted – which makes extensions more profitable than simple words. This game repeats for ten periods. In the presence of IP, the subjects are asked to assign a royalty fee to each of their creations. This royalty fee is paid to the word creator by any other player who might wish to extend the given word and is bound between zero and the value of the original word. In the absence of IP, the subjects do not pay any fee when extending other subjects' words.

The subjects play this game twice, with the same anonymous group. Each phase lasts for ten periods, which yields a total of twenty periods. Although an institution that is endogenously chosen by the players often produces a higher level of cooperation than an institution that is exogenously imposed on players, this difference is null or insignificant when the group has already played the game under an exogenous rule (Gallier, 2020). In our design, therefore, the intrinsic impact of democracy is minimized.

The subjects start in a baseline setting with IP and then move to a second unannounced phase in which IP has been abolished. The treatments impact the information that the subjects have at their disposal when transitioning and whether the subjects get to veto or not the institutional change by unanimous vote.

In the remainder of this section, the game and treatments are described in detail, and the main theoretical properties of the design are identified. The English translation of the original French instructions is available in Appendix A. Two videos that show the live play from our software interface are available at here.

3.1. Experimental task details

Preliminary control task

Before starting the main task, we ran a control task to measure the subjects' word-creation skills. The control task is built on Eckartz et al. (2012). All subjects are endowed with the same alphabetically ordered set of 9 letters (AMADITERS) and have 3 min to build as many words as possible by using only letters from the set that they can reuse as many times as they wish. The words must be at least 3 letters long.

Initialization stage

To jump-start the creation process, in both repetitions of the game, the subjects start the game with an endowment of five randomly drawn letters and two pre-created words. These words are associated with royalty fees of 50%. That is, each player starts the game as if they had already made two words for which they were charging 50% royalty fees. We chose royalty fees of 50% by default because this was the average value obtained in the first five periods in Brüggemann et al. (2016) (exactly 49%) of the IP treatment. This allows us to speed up the game and to offer equal extension possibilities to everyone at the beginning. In the real world, this corresponds to the public domain from which you

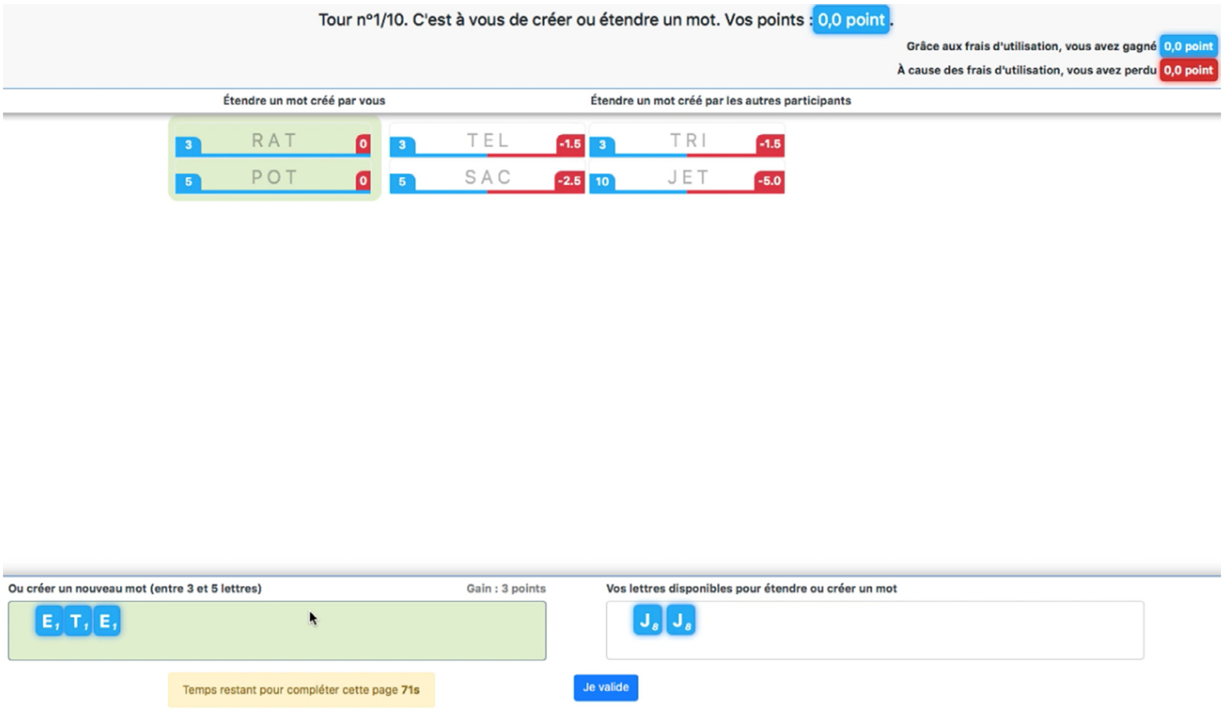

Fig. 1. General view of the main board at the beginning of the game.

can pick and that you can extend. The preselected words have been carefully chosen to have the same value and a set of similar extension possibilities. The same 6 words for the first part were given randomly to the 3 players (2 each) for both the first and second parts so that all groups faced the same game board in each treatment. That is, the initial endowment is the same for all groups, although it may vary across individuals.

Creation stage

The subjects play in turns. When it is their turn, the subjects can choose one among three possible actions:

1. produce a three- to five-letter word (*root*), for instance, if a subject has the letters a, e, i, m, g, and p, then she can create, among others, the roots *age, page* or *image*;
2. extend an existing word (*extension*). A subject is only allowed to extend an existing word with one additional letter, which can be placed in any position within the existing word. For example, *cat* can be extended, among others, to *cart* or *chat*; or
3. pass without taking action.

The subjects cannot recreate an already existing word. All produced words yield a payoff equal to the sum of the letter values (token), and 1 token converts to € 0.25. When extending a word, not only the one letter added but also all letters of the new word generate a payoff. For instance, extending *cat* (worth 3 + 1 + 1 = 5 tokens) to *chat* with the letter *h*, which itself is worth 4 tokens, results in a payoff of 9 tokens. This implements increasing returns in the experiment – whereby each subsequent innovation is *more* valuable than its components.[1]

A word is accepted when it is included in the French *Official Scrabble Dictionary 2018* implemented in the game. All forbidden actions are dynamically indicated to the player. If the word created or extended does not exist in the Scrabble dictionary or has already been created by another player, then a pop-up message warns the subject, who can correct her action in real time.

All the words and extensions created by each group member appear on the common board, as in Scrabble. For each word, the board shows its value, the amount of royalty fees to be paid to extend it and/or the royalty fees that are received in case another subject extends it. A screenshot of the board is given in Fig. 1.

[1] Although this assumption does not represent the structure of all industrial sectors, it is a faithful representation, for instance, of the software industry. See Brüggemann et al. (2016) for a discussion of this assumption.

IP stage

After having created a root or an extension, the subjects must set a royalty fee. A royalty fee for a root refers to all the added letters. A royalty fee for an extension refers only to the marginal contribution by the player – the one added letter. Royalty fees range from 0% to 100% in steps of 10%. The chosen royalty fee becomes public information and is fixed for the rest of the game. However, no one can be excluded from using the word altogether. By choosing higher royalty fees, the subjects earn more when their word creation is extended by other subjects. Nevertheless, at the group level, royalty fees are merely a mechanism to redistribute income within the group, as there are no transaction costs.

For example, Subject A produces *cat* – worth 5 tokens – and chooses a royalty fee of 50%. Every subject who extends *cat* pays A 2.5 tokens. Consider Subject B extending *cat* into *chat*, which is worth 9 tokens: 2.5 tokens are transferred to Subject A, Subject B earns 6.5 tokens and has to set a royalty fee for the letter *h*, which is worth 4 tokens. If she chooses 50%, then the next subject who adds a letter to the word *argent* will have to pay 2.5 tokens to Subject A for *cat* and 2 tokens (40% of 4 tokens, namely, the value of the letter *h*) to Subject B.

When it is not the subjects' turn, the main board of the game is shown so that players can follow the game but are not active.

Theoretical properties of the design

The modified Scrabble task transposes in the laboratory a sequential, cumulative innovation setting inspired by the model of Bessen and Meurer (2008). In particular, the game induces strict sequentiality, since each extension is created from an existing word by adding one and only one letter. This transposes the concept of innovation in the laboratory by letting the subjects be creative within a familiar space that is vast but intuitively searchable. This models the effort needed to invent in a natural way: generating complex ideas requires both effort and the ability to stand on the shoulders of giants. Moreover, the space is countable, as it resides totally within the chosen dictionary. We can compute the number and value of all possible extensions – we call this the *extendibility potential index* of a word – and have a precise and complete map of the innovation space that faces our subjects. An extension is worth its marginal contribution (the added letter) plus the value of the root that it is using. "*This payoff structure best describes the situation of basic science – in which the first contributions lay down the foundations, allowing subsequent contributions to carry most of the value – or in the software industry – in which modern software technologies are built on thousands of algorithms, hardware, drivers, etc. that have been accumulated over the years.*" (Brüggemann et al., 2016).

3.2. Treatments

Treatments impact the *transition* from the first to the second repetition of the game and focus on giving players the right to *veto* the institutional change through an unanimous vote and on the amount of information that the subjects have when voting.

First, we exogenously manipulate the presence or absence of a vote. In the vote treatment, after the first part with IP, the subjects are informed that they will have to vote to decide whether to keep the IP regime or to abolish it, i.e., that there will no longer be the possibility to set royalty fees for the words created and that all words will be publicly available at no extra fee. The voting rule is unanimity: the IP regime is kept only if everyone votes for it; otherwise, the regime changes to no IP. It only takes one subject voting to abolish the IP rule for all 3 members of the group to play no IP. To avoid impromptu choices, the subjects are forced to ponder the decision for at least 20 seconds. The vote is anonymous, and only the final result is known. At the end of the vote, the subjects are informed of whether they have won or lost the vote and which IP regime the group will play in the second part of the experiment.

We choose a unanimous vote to maximize the number of losers (the players who voted to keep IP but lost), whose behavior interests us the most. The vote procedure aims to implement in the laboratory a democratic process in which the subjects are forced to take a position and commit themselves to a clear choice. In addition, it simulates the disappointment that some would have if the abolition of IP were the result of a real political agenda. Most experiments rely on between-subjects designs in which each subject only experiences one treatment. However, policy changes in the real world never happen in a vacuum. Often, people start off in a new environment by knowing the old environment or people can experience only one system but know that another system is in place somewhere else. Previous experiences generate expectations and entitlements, and these might cause frustration from an imposed change. This frustration is illustrated, for instance, by the ultimatum game (Güth et al., 1982; Thaler, 1988) in which a mismatch between expectations and the decision of others usually leads individuals to retaliate against the decision maker, even if this retaliation is costly. In our case, this could lead the losers to exert less effort when they are forced by the result of the vote to play in a setting without royalties.

In its generic meaning, retaliation requires that subjects be ready to hurt themselves in order to hurt others. However, in our setting, this takes a specific form. The typical example of retaliation that we have in mind is *closing* rather than *opening* words. Closing words means, for example, adding an *s* after names or verbs. This generates a small payoff but is not a creative action that opens new possibilities for further extensions. Opening words means creating long words that allow others to have a large number of possible new combinations. In real life, this can be compared to applied vs. fundamental innovations that help others innovate. Of course, this is not necessarily a retaliation. Without IP, the incentives to create *opening* words are lower – there are no expected stream of royalties from further extensions – and this alone could explain why creators could switch to *closing* words, with no reference to retaliation. In many cases, it is difficult to evaluate what the

Table 1
Distribution of the subjects and groups across treatments.

Treatment	Info	Vote	Participants	Groups
Control	No	No	75	25
Control with information	Yes	No	75	25
Vote without information	No	Yes	150	50
Vote with information	Yes	Yes	165	55
All			465	155

best option is because, beyond the payoff itself, people calculate the opportunity to keep letters or not for the next move. When the word that maximizes points is unclear, we hypothesize that they prefer to minimize the advantage for others. In real life, this could mean that after a vote against IP, the research on applied (rather than fundamental) innovations increases more than what is predicted than if innovators optimized their research.

Second, we manipulate the information given to the subjects after the first part. In the information treatment, we inform the subjects just after the end of the first part that a previous experiment (Brüggemann et al., 2016) showed that the no IP regime significantly increases the total gain of subjects (see the instruction in Appendix A).

The degree of information provided to a voter has a significant impact on the final choice (Luskin et al., 2002). With this treatment, we want to simulate the effect of an informed vs. naive vote on the acceptance of the result of the vote by the losers. People can be disappointed for the wrong reasons, such as thinking that other people voted against them for their own selfish interest. However, if the result of the vote follows the recommendations of experts who provide evidence that the vote serves the common good, then disappointment should be lower.

We implement a situation with no vote, with or without information, as a control. The final breakdown of subjects by treatment is given in Table 1.

3.3. Experimental procedures

The experiments were conducted at the GAEL experimental laboratory (Grenoble, France). The experimental software was written in Python by using the oTree framework (Chen et al., 2016). The source code is available at here. We recruited 465 participants from the GAEL subject pool, which was composed of students and subjects taken from the general population in roughly equal parts. The subjects were allocated to 155 groups of 3 players. The instructions were read aloud and projected overhead. Clarification questions were answered collectively. The participants received a 10€ show-up fee, which was not at stake during the game. After the end of the experiment described in this paper, the subjects participated in an unrelated second experiment. Gains from the current experiment were used as an endowment for the second experiment, whose aims and results are described in Varaine et al. (2021). The participants knew that a second experiment would take place but were given no other details. After the second experiment and after being notified of their final payoff, the participants were asked to complete a short sociodemographic questionnaire, which also included two questions about whether the participants thought of themselves and the other players as being cooperative during the game.

The participants were on average 26 years old (std. dev. 10.2), and 58.5% were female. Sessions lasted approximately 90 min. The 465 participants earned 19.6€ on average, including a 10€ show-up fee, with a minimum payoff of 10.6€ and a maximum of 45.3€.

4. Results

4.1. Effects of abolishing IP

Given our focus on the effects of the transition to a no IP regime for society as a whole, all the results that we report in this section are generated from group-level data. Each group provides two observations, one per repetition. All results are robust to focusing on individual-level data.

Despite the differences in the details of our design (a within- vs. between-subjects design, 3 vs. 4 group size, 10 vs. 25 period, French vs. German language[2], starting endowment of words,...), we replicate the main result of Brüggemann et al. (2016): abolishing IP has a positive effect on innovation.

The 11 groups that play twice under IP show a small and not significant improvement in the second period compared to the first (+6.73 points, +5.04%, t test p value $= 0.553$, Cohen's $d = 0.27$), while the 144 groups that transitioned to a no IP setting, with or without voting, show a large and significant increase (+17.1 points, +11.6%, t test p value < 0.001, Cohen's $d = 0.66$).

[2] Differences in the structure of the language could allow for higher or lower extension potential and thus generate bias in the results. However, this is not the case: the "extendibility", i.e., the potential to generate more or less extensions of words in the two languages is rather similar. For details, see Appendix B.

Table 2
OLS estimates of group points by repetition, IP transition, and vote treatment

	Points	
	All groups	Control only
Repetition 1	131.27 *** (10.04)	130.86 *** (9.96)
Kept IP	-4.11 (7.89)	
Transition to noIP	22.05 ** (7.92)	9.36 ° (5.53)
Word skills	0.58 *** (0.14)	0.64 * (0.26)
N	310	100
R2	0.16	0.08
logLik	-1438.54	-472.28
AIC	2887.08	952.57

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; ° $p < 0.1$.

The above results mix together groups that transitioned to no IP by *fiat* and others that transitioned to no IP after a vote. The result might therefore be confounded by selection effects. A cleaner test of the hypothesis is to restrict attention only to the groups that did not undergo a vote and had no possibility to veto the transition. Restricting the sample to the 50 groups in the no-vote control treatment (with and without information), there is a weakly significant positive effect (+9.36 points, + 6.1%, *t* test p value 0.0843, Cohen's $d = 0.33$)

The effect survives when controlling for the average skills in the preliminary word task of the players of each group, with the same pattern: strong and highly significant for the entire data and weakly significant only for the control. The results of a difference-in-differences ordinary least squares (OLS) estimate of the total group points across the repetition and IP regime are reported in Table 2, where the first column reports all groups, and the second column reports only the control groups that did not vote.

These results hold true at the individual level. Transitioning to no IP significantly increases the individual earnings by 11.58% from 49.2 to 54.9 (+ 5.7 points, *t* test p value < 0.001), while sticking to IP results in a small, not significant 5.03% increase from 44.5 to 46.7 points (+2.24 points, *t* test p value $= 0.598$).

We accordingly replicate the general result of Brüggemann et al. (2016) of a positive effect of removing IP in a sequential innovation setting. Our results are weaker overall. This might be due to the far shorter length of the game (10 periods here, 25 in Brüggemann et al. (2016)), a factor that limits the sequences that could be generated by the players and, thus, the positive effect of the absence of IP on long sequences of innovations. The difference is unlikely to be due to language differences, however, as the languages do not differ much in terms of the *extendibility* potential – see Appendix B for details.

4.2. Effects of voting and information on the transition

In our preregistered analysis plan, we argued that the democratic nature of the transition would reduce the positive effect of transitioning to no IP. We hypothesized that the losers of the vote would feel acrimony toward the vote because it deprived them of royalty income and that they would react by reducing their efforts. This was just a dynamic version of the standard arguments in defense of IP: remove incentives to innovators, and they will exert lower effort.

Our hypothesis is not supported by the data. Actually, the results show the opposite pattern: voting has a *strongly positive* effect on the transition to no IP. To test this hypothesis, we compare the points accumulated by the groups that transitioned to no IP without a vote – i.e., all groups in the control treatments, irrespective of the amount of information received – to the groups that transitioned to no IP after a vote. We therefore compare 50 groups in the control treatments to 94 groups in the vote condition, irrespective of the information received.

After a vote, the groups that transitioned to no IP significantly and substantially increase their points (+21.23 points, +14.7%, *t* test p value < 0.001, Cohen's $d = 0.87$). This increase is weakly significantly different from the increase in the control treatments (*t* test, p value $= 0.0625$). This difference reaches traditional levels of significance when controlling for group word creation skills, as shown in Table 3, left column, which reports the results of a regression of the difference in the group points across repetitions by vote treatment, controlling for group skills.[3]

This result is robust to running the same tests at the individual level. The players in the control treatments increase their score by 3.12 points, while the subjects in the vote treatment increase it by 7.08 points. This difference is weakly significant in a direct test (*t* test, p value $= 0.0566$) and significant when controlling for skills (see Table 3, right column).

Overall, our preregistered hypothesis is at odds with our results. This result makes a test of the effect of information on the behavior of the losers of the vote irrelevant, and we thus skip this discussion altogether. Since information did have a negligible impact overall, in the remainder of the paper, we pool together the two information treatments and signal the results only when they do (slightly at best) differ.

[3] The regression here is carried out directly on the difference between the two repetitions rather than on difference-in-difference to make for a clearer reading. The results are of course identical to a difference-in-difference regression.

Table 3
OLS estimate of transition gains for groups transitioning to noIP, by voting condition.

	Transition gains	
	Group data	Individual data
No vote	5.48 (10.89)	2.66 (2.59)
Vote	12.16 * (6.10)	3.99 * (2.02)
Word skills	0.11 (0.27)	0.01 (0.06)
N	144	432
R2	0.03	0.01
logLik	-713.09	-1903.90
AIC	1434.18	3815.79

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; ° $p < 0.1$.

4.3. Determinants of the vote and the transition to no IP

Why did our hypothesis on the effects of voting fail? We argued that the best players would vote to keep IP, lose the vote, and react by reducing their effort. This did not happen. To understand what occurred, we venture here out of the preregistered plan and investigate individual voting determinants and their correlation with performance.

Let us start with a few descriptive statistics about voting. Without information, a slight majority of the subjects voted to keep IP (52%). Once given information about the results of Brüggemann et al. (2016), the majority favored a transition to no IP (64.8%). This difference is significant (Fisher exact test, p value < 0.001).

Given that unanimity was needed to keep IP, the vote resulted in a considerable number of losers – the subjects who voted to keep IP but were forced against their will into a game without IP. In the groups without information, there were 60 losers (out of 150 players), 20 groups had two losers, and 20 groups had one; among the groups with information, the number of losers decreased to 43 (out of 165), and 27 groups had one, and 8 groups had two.

The reasoning behind our hypothesis was simple. We assumed that the subjects' performance would be a function of skills and incentives. This means that under IP, the best players are rewarded for their production through royalties, and they therefore exert high effort, which results in increased productivity. The best players are the ones who earn the most royalties; accordingly, they will be the ones who vote against the transition to no IP. They are consequently the subjects most likely to lose the vote and feel robbed of a legitimate source of income. If they react by shirking, then the overall performance will suffer.

This argument can be split into a chain of hypotheses that can be tested with our data to shed light on why our hypothesis failed. In particular, we test whether a) skills in word creation correlate with output in the game, b) skills in word creation correlate with royalties earned, and c) royalties earned correlate with the likelihood of voting to keep the IP system in place.

Skills and performance

We hypothesized that skills would correlate with performance and royalty earnings. That is, the best players would end up performing better and being the winners of the IP system. Indeed, this is one of the reasons why proponents and advocates of IP support the existing system.

We measure skills with the results of the preliminary word-hunt task. This measure of skills is fully exogenous, as the subjects were not aware, when completing this task, of what they would be asked to do in the rest of the experiment.[4]

As expected, skills correlate significantly with the number of points accumulated (including net royalty gains) in the first repetition of the game ($r = 0.255$, p value < 0.001), in the second repetition with no IP ($r = 0.362$, p value $= 0.038$) and when sticking to IP ($r = 0.245$, p value < 0.001). In contrast, word-creation skills do not correlate with net royalty gains – the amount received because of royalties on created words minus the amount paid to access other participants' words. The correlation for the first repetition is small, negative in sign, and not significantly different from zero ($r = -0.05$, p value $= 0.279$). Word-creation skills also do not correlate with gross royalty inflow – which is a measure of the outreach of one's word on others without the impact of a player using other players' words for extensions ($r = -0.02$, p value $= 0.619$).

Extendibility potential

A further indicator of skills – and a way to assess how incentives shape the nature of play in our game – is to look at the *extendibility* potential of the roots created by the subjects. Words differ not only in terms of points generated upon creation but also in terms of how many and which extensions they can give rise to. IP rights could in principle incentivize the subjects to look for the most extendable words; these might yield few points upon creation but might generate sustained

[4] The results are robust to using an endogenous measure of skills, i.e., the performance of each subject relative to the potential choices that she could make at any point in time.

royalty streams as they allow the creation of more extensions. The "best" players could be those forward-looking subjects who create more flexible root words.

To examine this, using `python` and the official Scrabble dictionaries, we created a complete dataset of *all* the extensions that can be created from *any* 3-letter word (*root*) in French, German, and, for reference, English. For each root, we then counted the number of extensions and computed an *extendibility potential index*. We calculated the value of each extension generated from the root, and we deflated this value by the probability of being able to actually create the extension – that is, the probability of obtaining the letter(s) needed to create it.

Take as an example the word ACE (in English). From this word, 512 extensions can potentially be created – from FACE to SHACKLERS, HACKERS, and BRANCHES. Let us take FACE. Its value, i.e., the sum of the value of its letters, is 9. However, to create it, you need an F. The probability of drawing an F from the letter set in English Scrabble is 0.02. The deflated value of this word is thus $9 \times 0.02 = 0.18$. From FACE, you can build FARCE. To do this, however, you need *both* an F and an R – and the R has a probability of 0.06 of being drawn. FARCE is worth 10; its deflated value is $10 \times 0.02 \times 0.06 = 0.012$. Naturally, longer extensions are deflated more, as the likelihood of creating them decreases. Nonetheless, the *extendibility potential index* monotonically increases in the number of extensions that a root can generate and in the value of the letters added – which are essential properties.

In practice, however, the extendibility potential of words matters little. Although as expected and somewhat by construction, the index does significantly correlate with the amount of royalties received in repetition 1 ($r = 0.197$; p value < 0.001), it correlates with little else. The index does not correlate with the points earned for any repetition and treatment (Repetition 1: $r = 0.02$, p value $= 0.71$; repetition 2, IP: $r = -0.26$, p value $= 0.29$ repetition 2, no IP: $r = 0.04$, p value $= 0.36$). It unexpectedly does not significantly correlate with the inflows of royalties in repetition 2 among the subjects who chose to stick with IP ($r = 0.253$; p value $= 0.310$). Furthermore, the index does not differ by treatment (Repetition 2, mean index for IP: 2.79, for no IP: 2.68, t test, p value $= 0.70$) and is not different across the subjects who voted to keep IP (mean index 2.79) or to change to no IP (mean index 2.76, t test, p value $= 0.88$).

The index *does* make a difference in the reaction and composition of the subjects who lost the vote. The losers of the vote, i.e., the subjects who wanted to keep IP but were forced into a world with no IP, significantly reduced the extendibility potential across repetitions (from a mean index of 3.16 to 2.38, p value $= 0.002$, Cohen's d 0.505). This is the only evidence that we have of the losers reacting to losing the vote by changing strategy – possibly because of them being aggrieved. Nevertheless, these subjects were *not* the best players. The extendibility potential index in the first period correlates *negatively* with skills ($r = -0.122$, p value $= 0.032$).

Overall, despite some limited evidence that the losers of the vote *did* react by reducing efforts devoted to generating extendible root words, this had no bearing on the final result of the groups that transitioned to no IP, which performed better than under IP and far better than the groups that stuck with IP. The fact that extendibility played a minor role may be because it was perhaps a metric that was difficult to capture and that focusing on the extendibility of short words was something that *lower skilled* players would do – higher skilled players perhaps devoted their interest to longer words.

Skills and voting

Turning to voting behavior, skills show a weakly significant *negative* relation to voting for IP. IP voters had a mean skill of 107.99 points, while no IP voters had a mean skill of 119.89 (t test, p value $= 0.073$, Cohen's $d = 0.205$). The net royalty gains instead show a strong, significant and *positive* relation to voting to keep IP in place. IP voters had on average a net royalty balance of 3.77 points, which means that they received more royalties than they spent, while no IP voters had a negative net royalty balance of -2.86 points (t test, p value < 0.001, Cohen's $d = -0.75$).

The unexpected result of this analysis of the vote is therefore that the best players did not tend to lean toward keeping the IP system in place. Rather, the players who benefited from the IP system did vote to defend it, and they were not the best players.

Given that the unanimous voting rule imposed on several players who voted for IP a transition nonetheless to a no IP society, there is no significant impact on the average skill of the groups in repetition 2 (an average skill of 29.3 for the groups that kept IP in repetition 2 vs. 32.9 for the groups that transitioned to no IP, t test p value $= 0.222$, Cohen's $d = 0.21$).

Other determinants of voting and performance

However, skills can affect the individual impact of an IP system. In our game, there are at least four main drivers of net royalty earnings: skills, luck, greed and royalty aversion. High-skilled subjects can create complex and valuable words and thus gain royalties on them. Lucky subjects can reap the benefits of a better initial set of letters: not all letters have the same potential, and the set of letters was randomly drawn for each individual. Greedy subjects can set high royalty fees on their words to potentially obtain a larger stream of inbound royalties. Finally, if the subjects are averse to spend on royalties, then they can avoid extending other subjects' words and therefore reduce outbound royalties.

To make the analysis as easy to follow as possible, we create four variables that span the same domain for each of the abovementioned drivers of individual behavior (skills, luck, greed and royalty aversion).

For skills, we take the raw number of points generated in the preliminary word-creation task and rescale it so that the best player has a score of 100.

Table 4
Determinants of performance, royalty balance and voting.

	Points OLS	Royalties OUT	Royalties IN	Royalties BALANCE	Vote for IP	Vote for IP
constant	-10.214	36.003 ***	-2.219	-38.222 ***	0.117	-0.001
	(6.498)	(1.728)	(3.047)	(3.778)	(0.260)	(0.764)
skill	0.207 ***	0.041 ***	-0.023	-0.063 *		-0.005
	(0.045)	(0.012)	(0.021)	(0.026)		(0.005)
luck	0.340 ***	0.008	0.066 °	0.057		0.008
	(0.078)	(0.021)	(0.036)	(0.045)		(0.008)
greed	0.070 °	0.018 °	0.036 °	0.018		0.008 °
	(0.039)	(0.010)	(0.018)	(0.023)		(0.004)
royalty aversion	0.405 ***	-0.378 ***	0.075 **	0.453 ***		-0.005
	(0.061)	(0.016)	(0.029)	(0.036)		(0.008)
points					-0.006	-0.006
					(0.005)	(0.006)
royalty balance					0.056 ***	0.060 ***
					(0.009)	(0.011)
N	314	314	314	314	315	314
R2	0.230	0.642	0.044	0.352		
logLik	-1249.670	-833.689	-1011.865	-1079.350	-193.998	-189.651
AIC	2511.340	1679.378	2035.729	2170.699	393.997	393.303

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; ° $p < 0.1$.

We measure luck as the potential number of points that could be generated in period 1 for each subject. To do this, we generate all possible words and extensions that a subject could create given the available words and the allotted letter endowment at the start of the game, compute the value of each one of them, and take the mean. This number indicates the raw potential at the starting block for each player, irrespective of skills. It does show sufficient variance to in principle inhibit the performance in the rest of the game (mean 4.81 points, variance 1.07) – that is, despite our best efforts to level the playing field, one could get a lucky start in our game. We take this indicator and rescale it so that the luckiest subject has a score of 100.

We measure greed as the mean share of the royalty fees asked by each subject for each of their contributions. This number spans from 0 to 100.

We measure royalty aversion as the complement to 100 of the mean royalty fee accepted by each subject for extensions. Suppose that a subject extended words that originated only from herself. She did not pay any royalty fee, so her royalty aversion is 100. Suppose instead that a participant paid 50% royalties on one extension, 70% on a second one, and nothing on a third one. Her mean paid royalty is 40%, and her royalty aversion is 60. This number spans from 0 to 100.

To shed light on the three steps of our argument above, we run OLS regressions of the points earned, the royalties received and spent and their balance by using the four indicators described above, and we restrict our attention to the first repetition of the game, where the subjects played under IP. Since all four variables span from 0 to 100, the coefficients can be easily interpreted as the impact that increasing the variable by 1/100 would have on the dependent variable. The results are shown in Table 4, Columns 1 to 4.

The first column shows the impact of these four characteristics on the performance in the game, as measured by the number of points. Skills and luck contribute positively to success, which is intuitive. Royalty aversion has a similar positive role because it reduces outbound royalties. Overall, hence, our predictors all accordingly contribute in the hypothesized direction to success in the game.

The second, third and fourth columns show the impact on the royalty indicators (outbound, inbound, and balance). The signs go in the expected direction for greed (slightly increasing inbound royalties) and royalty aversion (heavily decreasing outbound royalties). We see no effect of luck across the board and, surprisingly, no effect of skills. If at all, skills *increase* the *out*bound royalties. Overall, this means that the winners of the royalty system are not the best players but the more 'autarkic' ones, i.e., those that restrict their attention to their own words only.

To shed light on the determinants of the vote, we gather all that we have learned from the above analysis, and we run a probit regression of voting for IP (Table 4, last two columns). First, we assess only the role of points – i.e., of bare performance – and of royalty balance – i.e., of performance at exploiting the IP system, and we drop all other potential drivers. The results clearly show that the royalty balance is strongly correlated with voting to keep the IP system, while performance is not strongly correlated with it. Then (last column), we add our individual controls – skills, luck, greed and royalty aversion. Again, the only strong predictor of voting for IP is receiving a net benefit from the IP system.

Our results show that although skilled players produce more points in the game and that the royalties earned predict voting patterns, skills in the game do not predict royalty earnings. Moreover, the subjects do not seem to be able to focus on more extendible words or if they are, this ability is not impacted by the IP system. The redistribution given by royalties is surprisingly orthogonal to the players' output and is even *negatively correlated* with skills in the game. Thus, the subjects who voted to keep IP in place and possibly grieved their loss were not the best players. The success in royalties is best predicted only by *royalty aversion*, an index that captures the tendency of a subject *not to pay* royalties and rely on her own words instead.

The IP system does not benefit the best players but the players who extract rent from the system. This explains both the failure of our preregistered hypothesis and the dire performance of the groups that stuck with IP in repetition 2. The best players were not among the losers of the vote and did not have a particular predilection for IP; that they ended up in an IP group after the vote was a sign of adverse selection, as it meant ending up in a group of rent-seekers bent on reducing royalty expenditure at the cost of less innovation.

Overall, we see that the IP system in our experiment does not reward the skills of the players but is linked only to royalty aversion. As a consequence, in our experiment, the subjects who benefitted from the IP system and therefore voted to keep it in place are the ones who *contributed relatively less* to overall creativity. It is no surprise then that in period 2, the groups that kept IP underperformed relative to the no IP groups. Although these results could be artefacts of our game, we feel that they could also serve as a truthful description of the IP system in the real world.

5. Discussion

We expected that losing the vote would reduce innovators' motivation to create, but this is not the case. We emphasize in the results section above that this is because the losers of the vote are not the most creative players. However, there is also at least one other competitive explanation: intrinsic motivation. The Scrabble-like task can be quite fun to play, and this might crowd out grievances to push even disappointed losers to exert effort. To the extent that the act of inventing also requires more laborious and tedious tasks than in our study, the external validity of our findings are somewhat limited.

We have reason to think that an intrinsically motivating task could be a good representation of several innovative sectors. Inventors often cite intrinsic motivations as their main driver. The *Time Invention Poll in Cooperation with* Qualcomm (2013), with 10,197 respondents worldwide, showed that "*the love of inventing*" is the main reason that respondents give when asked "*What do you think mainly motivates people to invent?*" (but with great differences among countries, e.g., 38% in the U.S., 65% in China, 34% in Germany, 53% in the U.K.)[5]. Using survey data for over 2000 academic scientists and engineers, Sauermann et al. (2010) found significant relationships between researchers motives and their patenting activities. Pecuniary motives predict patenting in the physical sciences, but the desire to contribute to society is the key motive that predicts patenting in the life sciences. In engineering, patenting is predicted by the motives of challenge and advancement. Although 74.5% of researchers rated the "*applicability of research*" as highly important, only 11.1% rated "*seeking IP rights*" similarly.

Because of all of these results, the fact that our Scrabble-like design game is fun is not a weakness, insofar as it can reproduce intrinsic motivations that are largely present in innovative sectors. However, it is true that the stakes are much higher in the real world and that the role played by intrinsic motivation differs widely among technology sectors. Thus, our experience can translate well the grievances that may occur in one sector and poorly those in another. For instance, the risk of stifling innovation with the abolition of IP could be high in technological areas where R&D investments are considerable but copying costs are quite low (e.g., pharmaceuticals), while in the field of computer or web technologies, this is much less the case. In future research, it would be interesting to couple the innovative nature of the Scrabble-like game with a more tedious, real-effort task to produce a stronger sense of grievances.

We wanted to test Machlup's point that institutional change in IP might be costly, and we focused on what we thought was a likely backlash from the best innovators being deprived of their royalties by the reform. Our hypothesis felt natural, and it is the one made by most if not all supporters of IP: touch it, and we will lose genius inventors. Our results show instead that the best innovators are not the players who earn more royalties. This is because the best players are "network" players, who earn a lot from royalties but also give a lot in royalties to others. This nullifies the beneficial effect of their creativity in an IP regime. In our experiment, the players who focused on gaming the system – collecting royalties while carefully avoiding paying them – are the only net winners of the system. Does this reflect reality? We think so. For publicly traded U.S. firms, Bessen et al. (2018) found that private costs had exceeded private rents since 1999–2000 and that the trend in costs was sharply higher due to a surge in the number of patent trolls, which contributed to the increase in the gap. The significant costs of competing can be a disincentive to innovation in an area where entry is very high (Shapiro, 2000; Lemley and Shapiro, 2006). Torrisi et al. (2016) showed that 67% of patent applications were filed to block other patents. These studies indicate that patents are mainly used strategically to block competitors and not as a means to reap a financial reward for an invention. Moreover, they found that most patents (38% in Europe, 36% in the U.S. and 46% in Japan) are never used, only approximately 6 patents out of every 100 are licensed, and only 4 out of 100 are sold. As in our game, inventors do not seem to be the most apt at commercializing their IP and benefit from royalties but tend rather to keep it for themselves and try to avoid paying royalties to others.

Aghion et al. (2018) evaluated the wage spinoffs in Finnish companies for the period of 1988–2012 when filing patents. They found that inventors received only 8% of the total private return, and the rest was shared by entrepreneurs (44%), blue-collar workers (26%) and white-collar workers (22%). Inventors earned on average a wage premium of 5% post invention. Bell et al. (2018) showed that a patent application marks the peak of a successful career In innovation rather than an event that produces high returns itself. Five years after a patent is granted, patent royalties account for less than 3% of income

[5] The other items are social good, making money, necessity, and fame.

even for inventors with highly cited patents. The authors also found that an inventor's salary was little related to granting a patent, except for the top 5% of patents cited. Therefore, it seems that our game reproduces the real balance in which innovators do not obtain as much positive profits on average from royalties.

In our game, the best players, who are not the ones who earn the most royalties, do not grant particular support to the patent system when voting. Is this case in the real world? It seems so, insofar as the role played by the formal patent system is minor for inventors. Wajsman and García-Valero (2017), based on data from the Eurostat Community Innovation Survey 2012, found that trade secrets are used by 52.3% of innovative companies, and patents are used only by 31.7% of them. This number is much higher for large companies: 69.1% use trade secrets, and 52.8% use patents, compared to 51.2% and 30.4%, respectively, for small and medium-sized enterprises (SMEs). They concluded that the use of trade secrecy to protect innovation is more important than the use of patents, especially among small SMEs. Levine and Sichelman (2018) showed that is also the case for startups.

For many companies, being the first to bring a product to the market and benefiting from this competitive advantage is better than patenting. This is the main idea behind the *competitive rent* models of Boldrin and Levine (2008b). As simply expressed by Leiponen and Byma (2009): "*If you cannot block, you better run*". In addition to the first-mover incentive, there is also a wide range of market incentives (market segmentation, speed of technological change, complementary sales, service and manufacturing capabilities, knowledge imperfections, transaction costs, learning and switching costs and network effects), social norms and psychological factors that can explain "Intellectual Production without Intellectual Property" (Dreyfuss, 2010).

Empirical studies that investigate how R&D investments respond to patent protection strengthening have shown no measurable increase in domestic R&D (Sakakibara and Branstetter, 1999; Qian, 2007), and Sweet and Eterovic (2019) showed no effect of IP on productivity in either developing or industrialized countries. These null results could be spurious, however, as the empirical studies of the effects of patent law changes face several limitations. Patent law changes may impact domestic R&D through channels that are difficult to track. For instance, through changes in foreign competition, changes in the foreign R&D investments of nondomestic firms that sell products to domestic consumers and, as noted by Budish et al. (2016), because technologies are developed for a global market, country-specific patent law changes in "small" economies may be a relatively small source of variation in global R&D incentives.

The literature on the "negative space" – creative and innovative fields that for historical, doctrinal, or other reasons, are not addressed by IP law (Raustiala and Sprigman, 2006) – is directly confronted with the widely held belief that copy protection is necessary for innovation incentives. Boldrin and Levine (2013) and Raustiala and Sprigman (2012) argue that in many industries, the first-mover advantage and the competitive rents that it induces are substantial without patents. The indirect incentives found in the open-source movement are a typical example (Lerner and Tirole, 2005). Research on intrinsic and extrinsic motivations (Benabou and Tirole, 2003) may also suggest that many managerial and institutional solutions are possible to compensate for the incentive deficit that would result from the removal of a temporary monopoly on inventions.

However, notably, some technological areas would probably suffer more than others from the removal of IP rights. In particular, the pharmaceutical and biotech industries would be strongly impacted (Sampat, 2018). Graham et al. (2009) asked through a large survey how important (or unimportant) the following seven items were to companies in securing competitive advantages from their technology innovations: first-mover advantage over competitors; secrecy; patents; copyrights; trademarks; difficulty of reverse engineering; and other production, implementation, or marketing capabilities. For medical device startups and venture-backed IT hardware companies, the respondents ranked patenting second behind the "first-mover advantage", but for biotechnology, patenting was ranked as the most important means of capturing competitive advantage.

6. Conclusion

In this paper, we focused on institutional change in the realm of IP rights. Through a Scrabble-like game, we recreated in the lab a sequential innovation setting and investigated the role of voting and information on the transition from a setting with IP to one in which IP has been abolished. We found that contrary to our expectations, the most creative players do not have a particular tendency to vote to defend IP; rather, the subjects who overwhelmingly voted to keep IP in place are those extracting rent from the IP system. Although not surprising *per se*, this finding is at odds with the traditional arguments put forth to keep IP in place, which mostly assume that the system obviously benefits the best innovators. This is not the case in our experiment.

The environment that we created is, however, hardly extendable to IP in general. We have mentioned the entertaining nature of the experiment: the subjects derive satisfaction from the game, especially those who are better at it. In this respect, this experiment can serve as a point of comparison for future research in environments where the creators would be less motivated to create. In addition, another important limitation must be pointed out. Even in the no IP experiments, it is only possible to freely capitalize on the work of others, but it is impossible to exactly copy it – if CAT has been produced, it cannot be produced again *verbatim*. Moreover, each created word yields earnings so that the creators capture a substantial amount of the returns from their innovation. This is not the case in many fields. Nordhaus (2004) estimates that only 2.2% of the social returns from technological advances are captured by the innovators, which is clearly lower than in our game. We have addressed this issue by leaving each creator to fix the value of their IP to ensure that the IP can produce more incentives. However, this has led to an increase in the price of royalties, has certainly reduced the incentives to use it and has inflated the number of 'autarkic' strategies.

Despite these limits, this experiment emphasizes that under specific environments that feature increasing returns and intrinsic motivations, patents can miss their target to benefit more rent seekers than from innovators. Even though it could be argued that these environments are rare in real life, some properties of the environment that we created are likely to exist in many fields. This can explain why, for example, patents are not the preferred incentive or protection system for innovators, except in the pharmaceutical industry (Sampat, 2018). Additionally, in our experiment, the players who voted to keep the IP system in place were on average less skilled than those who voted to abolish it. According to the *Time Invention Poll*, it is likely that the same would happen in the real world: support for the patent system is widespread in the population, but at the same time, the system is shunned, rarely used, and not truly valued by inventors. Among the general population worldwide, approximately 90% of the 10,197 respondents think that "patent protection promotes creativity and invention", and the same proportion think that "Patent systems are necessary to promote new invention and to guarantee that the inventor is given credit both intellectually and economically for his or her inventions". Only 10% think that the system "hurts" creativity and invention, and similarly, in another question, 10% think that "Patent systems are harmful because inventors are worried about being accused of infringing on the patents of others".

Moreover, this paper argues against the idea that institutional change is in itself detrimental. Contrary to what Fritz Machlup thought, reforming IP, in itself, is not necessarily costly for either innovators or the overall welfare. Our findings align with Buchanan and Wilson (2014), who used a different experimental design to model the effect of a change from an IP to a no IP treatment on innovators. They found no decrease in investment from innovators in the no IP treatment and eventually no decrease in overall welfare. Additionally, our findings suggest that in fields in which rent-seekers abound, the losers of reforms that democratically abolish IP are not innovators. Those who extract rents from the current IP protection system sometimes enjoy significant lobbying power. For example, it is very well-documented that Disney lobbied hard multiple times to extend the term of exclusivity offered by copyright when Mickey Mouse was about to enter the public domain (Landes and Posner, 2003) which is something no economist has ever supported, as prolonging an existing rent has no effect on further innovations.

As much as we believe in our results, we do not venture to provide direct policy advice. Our results have to be developed and clarified. According to the properties of each field, the policy of patents could prove to be somewhat effective. However, our experiment demonstrates that IP must not be considered an unconditional right to be constitutionally protected against a popular vote. Past scholarship has mainly defended IP constitutional protections against interest groups who over-represented rent-seekers' preferences (Nard and Morriss, 2006). However, these provisions generally produced qualified majorities that favored conservative minorities. These minorities – in the case of IPs – are mostly made up of rent-seeking users of IPs. Constitutional provisions are therefore likely to produce inefficiency. The problem of rent-seeking in patent systems invites a discussion of possible reforms that reduce it. Possibilities include making patent examination more stringent, ensuring that only patents that are valuable are granted, and more generally reforming IP to avoid patent trolls as much as possible. However, there is a debate as to whether it is worth spending more money to encourage the patent office to put more effort on examination (Lemley, 2000; Frakes and Wasserman, 2019). Another way is to implement a mechanism that reduces the risk of "submarine torpedoes", as Boldrin and Levine (2008a) call it, that is, the fear of being attacked by an unsuspected patent holder. Increasing the use of compulsory licenses or allowing independent invention claims could help in this case (Gallini and Scotchmer, 2002). Finally, more ambitious measures such as the patent buyouts mechanism imagined by Kremer (1998) could also help create an IP system that is advertised as highly as in the U.S. Constitution to serve innovators and "*To promote the Progress of Science and useful Arts*".

Declaration of competing interest

The authors declare to have no competing interests of any sort over this research. The financial support of the small grant call of INNOVACS is gratefully acknowledged.

Data availability

The full replication package (data and script) is available at the OSF page of the paper https://osf.io/uckbw/.

Acknowledgments

Financial support from INNOVACS is gratefully acknowledged.

Appendix A. Instructions for all treatments

The differences between treatments are indicated in square brackets. The original instructions were in French, participants also viewed screenshots and explanatory videos of the game. The original instructions are available from the authors upon request. The videos projected as a complement of the written instructions are available here.

Individual Training

To begin, before playing the paid game, you will do individual training. During this training, your task is to form as many words as possible in 3 minutes as in a Scrabble game. To do this, you have 9 letters at your disposal, which you can reuse as many times as you want. The words you form must be at least 3 letters long. All French words and their different declinations and conjugations are accepted. On the other hand, proper nouns, foreign language words, abbreviations and compound words are not accepted. Each time you find a word, you must enter it using the keyboard and press the “Enter” key on the keyboard to validate. If there is an accent in the word you wish to enter, you must type the letter without the accent.

Game, first repetition

In this part, you will be randomly divided into groups of three participants. Each in turn, with the other two participants in your group, you will have to form words, just like in the training task. You will play over ten periods. Each period corresponds to one round of the game for each of the participants in your group. When it is your turn to play, you have two options:

- Create a new word.
- Extend an existing word to create a new one.

You have 40 seconds to make your decisions. If you have not created a new word by the end of the 40 seconds, you pass your turn. To begin, you have five letters at your disposal, as well as two words already formed. Then you will receive a new letter for each new round.

Word creation

The first option is to create a new word. To create a new word, you have to drag the letters at your disposal with your mouse from the frame at the bottom right to the “Create a new word” frame at the bottom left of the screen. Once the new word has been created, you must validate it by clicking on “I validate”. Attention you must create words of at least three letters.

Word extension

The second option is to extend an existing word. You can extend the word of your choice, whether it is one of your words or a word from another participant. To extend a word, you must drag one of your letters with your mouse and add it to an existing word. You can place the letter anywhere in the word - at the beginning, middle or end of the word. However, you can’t rearrange the other letters in the word. When you expand a word, you can only add one letter. Once you have added the letter, you must validate by clicking on “I validate”. When you create or extend a word, the letters you have used are deleted from your stock. The game checks each word entered with a spell checker. If the word is incorrect, the word creation field turns red. There is no penalty for entering a word incorrectly, you simply have to enter a new word.

Payoff

As in the Scrabble game, each letter has a certain value in points. When you create or extend a word, you earn the sum of the points for the letters of the newly created word. For example, if you create the word $P_3O_1T_1$, you earn 5 points = 3 points for P, 1 point for O, 1 point for T. Then, if you expand the word $P_3O_1T_1$ into $P_3O_1R_1T_1$, you earn 6 points = 3 points for P, 1 point for O, 1 point for R and 1 point for T.

Royalties

One last rule that will determine your points: user fees. When you create or extend a word, you have to choose an amount of fees for the use of your letters by other players. This way, players who wish to extend your word will have to pay you a fee. When validating your word, you must choose the level of fees as a percentage of the value of the letters you contribute, from 0% to 100% depending on your choice. To use the example of $P_3O_1T_1$, if you choose a 50% user fee, then players who wish to extend the word will have to pay you 50% of the value of $P_3O_1T_1$, i.e. 2.5 points.

Earnings

So at the end of each period you earn points. You earn the value of the word you created and the fees received from other players, minus the points paid in fees to other players. At the end of the game, you will see the number of points you have accumulated as well as the total fees you have paid and received. Your final number of points will contribute to your final payment in Euros at the end of the session, at the conversion rate 1 POINT = 0.25 EUROS.

Game, second repetition

Vote/noVote

Now that you've completed the first part of the word game, we'll move on to the second paid part.

[noVote condition: But unlike the first game, from now on you will play without word fees: you won't pay any more fees to use other players' contributions and you won't earn any more points when other participants use your contributions.**]**

[Vote condition: Before playing, we will ask you to vote with the other participants in your group on the rules you prefer for this second game. You must indicate your preference between these two options:

- Remove word fees from the rules of the game: you will no longer pay any fees to use other players' contributions and you will not earn any more points when other participants use your contributions.
- Keep the fees for using words: the same rules as the previous game will apply.

You will play a game with user fees if all participants in your group vote to keep the fees. Otherwise, you will play a game without user fees. Before voting, you will have 20 seconds to think about your choice.**]**

Information

[Info condition: For your information, a scientific study published in 2016 in the journal Research Policy analyzed the behavior of participants in a game similar to the one you are playing. The study compared the performance of participants in a game with the word user fee rule to a game without the user fee rule. The findings of the study clearly show that, overall, participants performed better in the game without a word user fee than in the game with a user fee. In a game without user fees, participants overall produced more word extension and more points. As a result, this resulted in higher payments for players in the no-cost game.**]**

Appendix B. Extendibility across languages

Languages can and do differ in many ways. Related to our game, languages could offer different levels of flexibility for creating extensions; that is, they can be more or less "dense" in the number of words that you can create starting from a 3-letter word and adding one letter in any position – our rules for extensions. This is important when comparing our results – obtained in a French-speaking setting – with the results of Brüggemann et al. (2016) – obtained in a German-speaking setting. Indeed, the German vocabulary *is* larger than the French vocabulary by a wide margin.

We built the *extendibility potential index* for any 3-letter word in English, French and German. We then analyzed whether we could see any systematic differences across languages. The answer is *no*.

To be more specific, the distribution of French and German words is not significantly different in the mean (*t* test, p value = 0.07), distribution (Kolmogorov-Smirnov test, p- value = 0.156), or nonparametric test, which better reflects the nonnormal nature of the data (Mann-Witney test, p value = 0.534).

Moreover, the distributions do not show any striking difference, as seen in the plot below, by plotting both the pdf and cdf of the *extendibility potential index* for the three languages.

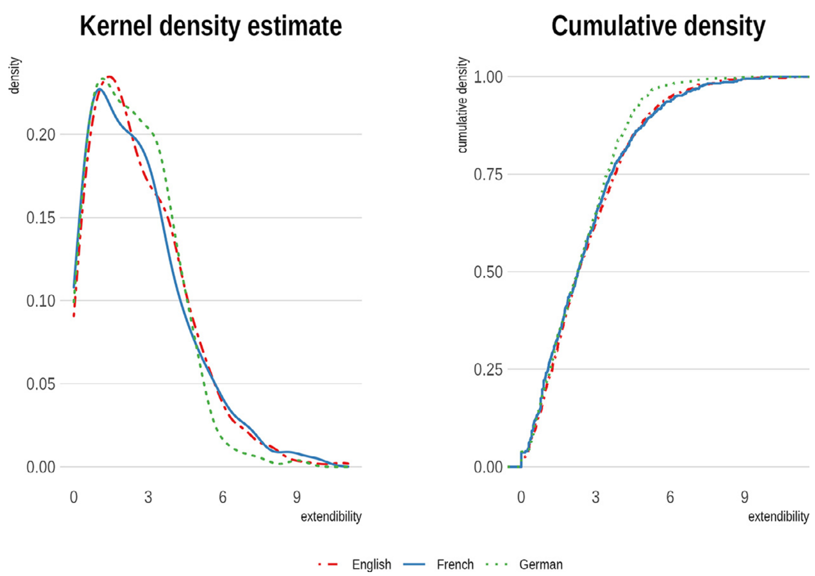

Fig. 2. Distribution of the extendibility potential index by language.

The reason why this index does not vary much with language is that Scrabble is a remarkably well-crafted game. The value of the letters and their frequency in the letter set in Scabble vary across languages *specifically* to reflect the relative scarcity/abundance and ease of use in each language. Our manipulation – adding one letter anywhere to the word – is different from standard Scrabble rules but not dramatically different; as a consequence, it benefits from the robust conception and development of the Scrabble letter values and scarcity.

We can then safely say that the differences in the results of the German and French experiments are *not* ascribed to differences in the underlying language structures in our game.

References

Aghion, P., Akcigit, U., Hyytinen, A., Toivanen, O., 2018. On the Returns to Invention within Firms: Evidence from Finland. In: AEA papers and proceedings, pp. 208–212.
Anderson, T., Hill, P.J., 1988. Constitutional constraints, entrepreneurship and the evolution of property rights. JD Gwartney and RE Wagner (1988), Public choice and constitutional economics 207–208.
Ariely, D., Gneezy, U., Loewenstein, G., Mazar, N., 2009. Large stakes and big mistakes. Rev. Econ. Stud. 76, 451–469.
Bell, A., Chetty, R., Jaravel, X., Petkova, N., Reenen, J.V., 2018. Who becomes an inventor in america? the importance of exposure to innovation. Q. J. Econ. 134, 647–713. doi:10.1093/qje/qjy028.
Benabou, R., Tirole, J., 2003. Intrinsic and extrinsic motivation. Rev. Econ. Stud. 70, 489–520.
Bessen, J., Meurer, M.J., 2008. Patent failure: How judges, bureaucrats, and lawyers put innovators at risk. Princeton University Press.
Bessen, J., Neuhäusler, P., Turner, J.L., Williams, J., 2018. Trends in private patent costs and rents for publicly-traded united states firms. Int. Rev. Law Econ. 56, 53–69.
Boldrin, M., Levine, D.K., 2008a. Against intellectual monopoly. Cambridge University Press Cambridge.
Boldrin, M., Levine, D.K., 2008b. Perfectly competitive innovation. J. Monet. Econ. 55, 435–453.
Boldrin, M., Levine, D.K., 2013. The case against patents. J. Econ. Perspect. 27, 3–22.
Brüggemann, J., Crosetto, P., Meub, L., Bizer, K., 2016. Intellectual property rights hinder sequential innovation. experimental evidence. Res. Policy 45, 2054–2068.
Buccafusco, C., Sprigman, C., 2010. Valuing intellectual property: an experiment. Cornell L. Rev. 96, 1.
Buccafusco, C., Sprigman, C.J., 2019. Experiments in Intellectual Property. Research handbook on the economics of intellectual property law. Edward Elgar Publishing.
Buchanan, J.A., Wilson, B.J., 2014. An experiment on protecting intellectual property. Exp. Econ. 17, 691–716.
Budish, E., Roin, B.N., Williams, H., 2016. Patents and research investments: assessing the empirical evidence. Am. Econ. Rev. 106, 183–187.
Charness, G., Grieco, D., 2014. Creativity and financial incentives. University of California, Santa Barbara, Working Paper.
Chen, D.L., Schonger, M., Wickens, C., 2016. Otree–an open-source platform for laboratory, online, and field experiments. J. Behav. Exp. Finance 9, 88–97.
Crosetto, P., 2010. To patent or not to patent: A pilot experiment on incentives to copyright in a sequential innovation setting. In: IFIP International Conference on Open Source Systems. Springer.
Dannenberg, A., Gallier, C., 2019. The choice of institutions to solve cooperation problems: a survey of experimental research. Exp. Econ. 1–34.
Dreyfuss, R.C., 2010. Does ip need ip-accommodating intellectual production outside the intellectual property paradigm. Cardozo L. Rev. 31, 1437.
Eckartz, K., Kirchkamp, O., & Schunk, D. (2012). How do incentives affect creativity?
Eidelman, S., Crandall, C.S., 2012. Bias in favor of the status quo. Soc. Personal. Psychol. Compass 6, 270–281.
Eidelman, S., Crandall, C.S., Pattershall, J., 2009. The existence bias. J. Pers. Soc. Psychol. 97, 765–775. doi:10.1037/a0017058.
Eidelman, S., Pattershall, J., Crandall, C.S., 2010. Longer is better. J. Exp. Soc. Psychol. 46, 993–998. doi:10.1016/j.jesp.2010.07.008.
Eisenberger, R., Rhoades, L., 2001. Incremental effects of reward on creativity. J. Pers. Soc. Psychol. 81, 728.
Eisenberger, R., Shanock, L., 2003. Rewards, intrinsic motivation, and creativity: a case study of conceptual and methodological isolation. Creat. Res. J. 15, 121–130.
Erat, S., Gneezy, U., 2016. Incentives for creativity. Exp. Econ. 19, 269–280.
Frakes, M.D., Wasserman, M.F., 2019. Irrational ignorance at the patent office. Vand. L. Rev. 72, 975.
Gallier, C., 2020. Democracy and compliance in public goods games. Eur. Econ. Rev. 121, 103346.
Gallini, N., Scotchmer, S., 2002. Intellectual property: when is it the best incentive system? Innovat. Policy Econ. 2, 51–77.
Graham, S.J., Merges, R.P., Samuelson, P., Sichelman, T., 2009. High technology entrepreneurs and the patent system: results of the 2008 berkeley patent survey. Berkel. Technol. Law J. 24, 1255.
Green, J.R., Scotchmer, S., 1995. On the division of profit in sequential innovation. Rand. J. Econ. 20–33.
Güth, W., Kocher, M.G., 2014. More than thirty years of ultimatum bargaining experiments: motives, variations, and a survey of the recent literature. J. Econ. Behav. Org. 108, 396–409.
Güth, W., Schmittberger, R., Schwarze, B., 1982. An experimental analysis of ultimatum bargaining. J. Econ. Behav. Org. 3, 367–388.
Kaufman, J.C., Sternberg, R.J., 2019. The cambridge handbook of creativity. Cambridge University Press.
Kremer, M., 1998. Patent buyouts: a mechanism for encouraging innovation. Q. J. Econ. 113, 1137–1167.
Landes, W.M., Posner, R.A., 2003. Indefinitely renewable copyright. Univ. Chicago Law Rev. 70, 471–518.
Laske, K., & Schroeder, M. (2017). Quantity, quality and originality: The effects of incentives on creativity.
Leiponen, A., Byma, J., 2009. If you cannot block, you better run: small firms, cooperative innovation, and appropriation strategies. Res. Policy 38, 1478–1488.
Lemley, M.A., 2000. Rational ignorance at the patent office. Nw. UL Rev. 95, 1495.
Lemley, M.A., Shapiro, C., 2006. Patent holdup and royalty stacking. Tex. L. Rev. 85, 1991.
Lerner, J., Tirole, J., 2005. The economics of technology sharing: open source and beyond. J. Econ. Perspect. 19, 99–120.
Levine, D.S., Sichelman, T., 2018. Why do startups use trade secrets. Notre Dame L. Rev. 94, 751.
Luskin, R.C., Fishkin, J.S., Jowell, R., 2002. Considered opinions: deliberative polling in britain. Br. J. Polit. Sci. 32, 455–487.
Machlup, F., 1958. An economic review of the patent system. US Government Printing Office, p. 15.
Meloso, D., Copic, J., Bossaerts, P., 2009. Promoting intellectual discovery: patents versus markets. Science 323, 1335–1339.
Moschini, G., Yerokhin, O., 2008. Patents, research exemption, and the incentive for sequential innovation. J. Econ. Manag. Strateg. 17, 379–412.
Moshinsky, A., Bar-Hillel, M., 2010. Loss aversion and status quo label bias. Soc. Cogn. 28, 191–204. doi:10.1521/soco.2010.28.2.191.
Nard, C.A., Morriss, A.P., 2006. Constitutionalizing patents: from venice to philadelphia. Rev. Law Econ. 2, 223–321.
Nordhaus, W. D. (2004). Schumpeterian profits in the american economy: Theory and measurement.
Ouellette, L.L., 2015. Patent experimentalism. Va. L. Rev. 101, 65.
Qian, Y., 2007. Do national patent laws stimulate domestic innovation in a global patenting environment? a cross-country analysis of pharmaceutical patent protection, 1978–2002. Rev. Econ. Stat. 89, 436–453.
Qualcomm, 2013. Time invention poll in cooperation with qualcomm. Time.
Raustiala, K., Sprigman, C., 2006. The piracy and paradox: innovation and intellectual property in fashion design. Va. L. Rev. 92, 1687.
Raustiala, K., Sprigman, C., 2012. The knockoff economy: How imitation sparks innovation. Oxford University Press.

Sakakibara, M., & Branstetter, L. (1999). Do stronger patents induce more innovation? evidence from the 1988 japanese patent law reforms. Technical Report, National bureau of economic research.
Sampat, B. N. (2018). A survey of empirical evidence on patents and innovation. Technical Report, National Bureau of Economic Research.
Samuelson, W., Zeckhauser, R., 1988. Status quo bias in decision making. J. Risk Uncertain 1, 7–59. doi:10.1007/BF00055564.
Sauermann, H., Cohen, W., Stephan, P., 2010. Doing Well or Doing Good? the Motives, Incentives and Commercial Activities of Academic Scientists and Engineers. In: DRUID summer conference, pp. 16–18.
Scotchmer, S., 1991. Standing on the shoulders of giants: cumulative research and the patent law. J. Econ. Perspect. 5, 29–41.
Shapiro, C., 2000. Navigating the patent thicket: cross licenses, patent pools, and standard setting. Innovat. Policy Econ. 1, 119–150.
Sprigman, C.J., Buccafusco, C., Burns, Z., 2013. What'S a name worth: experimental tests of the value of attribution in intellectual property. BUL Rev. 93, 1389.
Stiglitz, J.E., 2008. Economic foundations of intellectual property rights. Duke LJ 57, 1693.
Sweet, C., Eterovic, D., 2019. Do patent rights matter? 40 years of innovation, complexity and productivity. World Dev. 115, 78–93.
Thaler, R.H., 1988. Anomalies: the ultimatum game. J. Econ. Perspect. 2, 195–206.
Torrisi, S., Gambardella, A., Giuri, P., Harhoff, D., Hoisl, K., Mariani, M., 2016. Used, blocking and sleeping patents: empirical evidence from a large-scale inventor survey. Res. Policy 45, 1374–1385.
Tversky, A., Shafir, E., 1992. Choice under conflict: the dynamics of deferred decision. Psychol. Sci. 3, 358–361. doi:10.1111/j.1467-9280.1992.tb00047.x.
Varaine, S., Benslimane, I., Magni Berton, R., Crosetto, P., 2021. Attacking the weak or the strong? an experiment on the targets of parochial altruism. Political Behavior, online first.
Wajsman, N., García-Valero, F., 2017. Protecting innovation through trade secrets and patents: determinants for european union firms. Eur. Union Intellect. Property Office.