

Designing Open Source Licenses*

Kim-Sau Chung[†] Melody Lo[‡]

December 29, 2023

Abstract

Open source licenses are noted for being self-referential. The two most popular licenses are GPL and BSD. GPL says the next developer cannot go proprietary, and can only go open source with the same license, namely GPL. BSD says the next developer can go proprietary, and can also go open source with any license, including BSD. We provide a framework to study all self-referential licenses, including but not limited to GPL and BSD. We provide an explanation of why GPL and BSD stood out from all self-referential licenses as being the most popular among open source developers. We also present settings where there are other self-referential licenses better than GPL and BSD.

KEYWORDS: open source, licenses, self-referential, GPL, BSD

JEL CLASSIFICATIONS: L86, O36

*We thank Jimmy Chan, Yeon-Koo Che, Yi-Chun Chen, Sambuddha Gosh, Junichiro Ishida, Chia-hui Lu, Thomas Mariotti, Balazs Szentes, William Taysom, Kelvin Yuen, and participants at various seminars and conferences for very helpful discussions. Chung and Lo are co-first authors. Chung acknowledges the support of the Research Grants Council of Hong Kong (via grant GRF12502121) and HKBU (via grant RC-IG-FNRA/17-18/01). Lo acknowledges the support of the Higher Education Sprout Project of the Ministry of Education in Taiwan (via grant 113L900303 to the Taiwan Social Resilience Research Center).

[†]Centre for Business Analytics and the Digital Economy, Hong Kong Baptist University; email: kimsauchung@gmail.com

[‡]Department of Economics, National Taiwan University; email: peiyulo2006@gmail.com

1 Introduction

Once a software is developed, its developer has at least two ways to distribute it. The first is to go proprietary, meaning that she sells copies of the binary code for a profit. Since the binary code is difficult to interpret, it deters the others from disabling its security device and making illegal copies. But it also makes it difficult for future developers to learn from her source code. In other words, while going proprietary allows the developer to profit from her effort, it also stifles further development of her original idea.

Another way the developer can distribute her software is to go open source, meaning that she makes her source code open for everyone to see and copy. By going open source, she forgoes the profit she could have made from selling copies of the binary code, but she can gain utilities when future developers, inspired by her source code, develop more advanced versions. Such utilities may come from pride, prestige, sheer joy of seeing her idea further developed, or the satisfaction of using a more advanced version of her original software.

There are two major kinds of license a developer may use when she decides to go open source. The first is the restrictive kind, exemplified by the GPL license.¹ By sharing her source code using GPL, the developer tells future developers, “You can use my source code to develop a more advanced version. But if you ever want to distribute your advanced version, you have to go open source instead of going proprietary, and you have to go open source with the same license I am using, namely GPL.” This kind of open source licenses are restrictive because they restrict how future developers can distribute their softwares. In particular, the restriction is a “share-alike” requirement, requiring that future developers share in the same manner as the original developer.

One of the most famous developers who went open source using GPL is Linus Torvalds. Ever since Linus Torvalds distributed the source code of the first version of Linux using GPL, all subsequent versions have to be shared alike using the same

¹GPL stands for “general public license”.

license. Even when Red Hat, a for-profit company, developed their commercial version of Linux, called Red Hat Enterprise Linux, they were also bound by GPL to share their source code. It means that they could not make a profit by selling copies of the binary code. Instead, they could only make a profit by selling complementary services. Today, Linux powers millions of smart devices, including smart phones, smart TVs, tablet computers, etc.

The second kind of license a developer can use when she goes open source is the permissive kind, exemplified by the BSD license.² By sharing her source code using BSD, the developer tells future developers, “You can use my source code to develop a more advanced version, and you can distribute your advanced version in whichever way you like. You can go proprietary, you can also go open source. If you choose to go open source, you can use any license you like, including BSD.”

One of the most famous developers who went open source using a BSD-like permissive license is Donald Knuth. While the core of \TeX is copyrighted by Donald Knuth and no changes are permitted, add-on programs are licensed under the \LaTeX Project Public License (LPPL), which is very much like BSD (Gaudeul, 2007). The permissive nature of LPPL allowed subsequent developers to go proprietary and profit from their efforts. This option encouraged many developers to participate in developing various add-on programs, with Scientific Workplace being one of the most profitable examples.

The open source movement is nothing short of a revolution in how production is organized. Many of the most valuable softwares (such as Linux, \LaTeX , Apache, etc.) might never have been developed if developers had not learned how to share their contributions using open source licenses. The movement also had impacts reaching far beyond the software industry. Projects with user-generated contents such as Wikipedia might never have been possible if contributors had not learned

²BSD stands for “Berkeley software distribution”. We use BSD as an umbrella term referring to various BSD-like licenses including, for example, LPPL (to be discussed below), MIT (which is even more permissive than BSD by not requiring acknowledgement of the licensor), and Apache (which makes explicit certain permissions that are only implicit in BSD). See, for example, Smith (2022).

how to share their contributions using Creative Common licenses, which in turn were inspired by open source licenses.

While GPL and BSD are the two dominant open source licenses, with their dominance especially pronounced at the early stage of the open source movement,³ it is important to recognize that they were inventions of idealistic developers, who invented them more as their anti-capitalist manifestos, instead of as calculated designs that maximize productivity. Therefore, it is conceivable that a more careful design exercise can unleash even more productivity.

A prerequisite of this design exercise is, of course, a theoretical framework to study all possible open source licenses, which is the first contribution of this paper. Since open source licenses restrict future developers' choices of open source licenses, they are necessarily self-referential in nature. Open source licenses hence resemble types in epistemic game theory, where a type describes a belief over types. A key element of our theoretical framework is hence the spaces of open source licenses, which have a similar structure as type spaces.

Our theoretical framework makes it easy to conceive of many more open source licenses other than GPL and BSD. Many of them cannot be easily pigeonholed into either the restrictive or the permissive camp. For example, a license may be restrictive in the sense that it restricts what restrictions the next developer can impose on the next-next developer, but is permissive at the same time exactly because the next-next developer is less restricted.

Once we can conceive of many more open source licenses other than GPL and BSD, a natural question is why they are not as popular among open source developers. One possibility is that the developers were simply not imaginative enough. Had they been equipped with our theoretical framework, they would have expanded their imagination and conceived of many more open source li-

³Vendome *et al.* (2017) study 16,221 Java projects on GitHub. They find that, by 2012, BSD-like licenses (including MIT and Apache) and various versions of GPL still accounted for more than 90% of all open source licenses used. Balter (2015) reports that, up to 2015, only 15% of open source projects on GitHub used a license other than MIT, Apache, or GPL.

censes other than the two everyone is using. Another possibility is that there exist simple, easy-to-understand models—perhaps not realistic but still good enough approximations of the reality—where GPL and BSD are indeed the two best open source licenses. These models are what developers subconsciously use to reason about their own community, and hence also lead them to stick with GPL and BSD.

We provide supports for both possibilities. In Section 6, we present a setting that is a plausible approximation of reality, where it is indeed the case that all other open source licenses can be ignored without loss of generality. We develop our argument in two steps. First, we introduce an axiom called imposture-proofness to exclude open source licenses that can be gamed by a developer by splitting her version of the software into two consecutive versions, with only the first version subject to the restrictions contained in the previous developer’s open source license. Among other implications of imposture-proofness, we show that any imposture-proof open source license that does not allow the next developer to go proprietary is essentially the same as GPL.

In the second step, we show that, with a linear structure of developers and exponential discounting, a developer going open source cannot do better than using either GPL or BSD. If she does not allow the next developer to go proprietary, then her license is essentially the same as GPL, per the result in the first step. If she does, then BSD gives the next developer maximum freedom, which is good for her as well. The key in this last step is that exponential discounting and a linear structure of developers render the interests of the current and the next developer perfectly aligned when it comes to putting restrictions on the next-next developer.

We also provide support for the other possibility. In Section 7, we demonstrate how the sufficiency result in Section 6 breaks down if we relax either the assumption of a linear structure of developers or the assumption of exponential discounting.

This paper is structured as follows. The rest of this section reviews the related literature. In Section 2, we start with a basic model that only allows for GPL

and BSD. In Section 3, we slightly extend our basic model by introducing three new open source licenses. This exercise illustrates how our theoretical framework allows us to conceive of new, never-heard-of open source licenses. The three new open source licences introduced in this section will also play important roles in our subsequent analysis. This section also explains why it can be difficult to summarize an open source license by a single parameter called “permissiveness” as in Lerner and Tirole (2005b).

In Sections 4 and 5, we introduce two important concepts, irreducibility and imposture-proofness, that are crucial in our statement of the sufficiency result, which in turn is presented and proved in Section 6. The sufficiency result says that it is without loss of generality to focus only on GPL and BSD under the assumptions of a linear structure of developers and exponential discounting. These two assumptions will be relaxed in Section 7, where we show how two of the new open source licenses introduced in Section 3 can each be better than GPL and BSD if one of these two assumptions are relaxed. Section 8 concludes.

1.1 Related Literature

The open source revolution has rightfully attracted a lot of economic studies. Lerner and Tirole (2002, 2005a) provide some early introduction of this revolution, and Fershtman and Gandal (2011) a brief survey of the related economics, to economists. Subsequent studies can be roughly divided into four strands. The first focuses on the competition between commercial software companies and the open source community—what prices the companies would set, the chance that they can survive this competition, etc. The open source community is typically modelled as a group of altruistic members who both contribute to and benefit from the open source movement. How their collaboration is facilitated or jeopardised by different choices of licenses is typically not the focus.⁴ Factors other than the

⁴Presumably the license being used is not BSD, as members do not have the option of going proprietary.

license are instead the focus. For example, Johnson (2002) focuses on the size of the open source community, Casadesus-Masanell and Ghgemawat (2006) on demand-side learning, and Economides and Katsamakas (2006) on network externalities.

This paper differs from this strand of studies in that it focuses explicitly on open source *licenses*—what options other than GPL and BSD do we have, can they serve purposes that neither GPL nor BSD can serve, and when can they be ignored without loss of generality?

The second strand is concerned with the kind of altruism that motivates members of the open source community—are they motivated by warm glow, pride, consumer surplus, or are they merely motivated by the material payoffs from signalling their competence? Athey and Ellison (2014) theoretically study how different kinds of altruism affect the competition between commercial software companies and the open source community; while Hertel, Niedner, and Herrmann (2003), Lerner, Pathak, and Tirole (2006), Roberts, Hann, and Slaughter (2006), and Fershtman and Gandal (2007) empirically study related questions.

In comparison to this strand of studies, our model makes specific assumptions on the kind of altruism that motivates a developer to go open source. In particular, we assume that he internalizes part of the consumer surplus that may be generated by future developers.⁵ We have not explored the implications of other kinds of altruism such as warm glow.

The third strand is interested in the governance structure of the open source community. Most of these studies are empirical in nature, taking the form of meticulous case studies of selected open source projects, and taking advantage of the fact that, for most such projects, the whole history of interactions is stored as log files in the public domain. Examples of these studies include Fielding (1999) on the Apache project, Mockus, Fielding, and Herbsleb (2002) on the Apache and Mozilla projects, and Han and Xu (2019) on the Python project. See also von Hippel and von Krogh (2003) and Johnson (2006) for some general insights distilled from

⁵Or, alternatively, he gains an “egoboo” everytime a more advanced version of his software is developed. See Footnote 8 for this this alternative interpretation.

these studies.

In comparison to this strand of studies, our model assumes that each generation of the software is developed by one and only one developer. The very interesting topic of intra-generation coordination hence cannot be studied in our stylized model.

The final strand is the only strand that explicitly studies open source *licenses*. It is also by far the smallest strand, with literally fewer than a handful of studies. Lerner and Tirole (2005b) provide a stylized model that differentiates different open source licenses by a single parameter, namely their “permissiveness”. Their model abstracts away the defining features of different open source licenses, such as GPL’s “share-alike” requirement and BSD’s permission to go proprietary, rendering it impossible to discuss the optimality of mixing-and-matching these features. As we shall explain in Section 3, “permissiveness” is also an ambiguous concept. One can easily construct examples that are permissive and restrictive at the same time.

In two unpublished working papers, Gaudeul (2004, 2005) compares GPL and BSD using a model less stylized than that of Lerner and Tirole (2005b). In particular, she explicitly models GPL’s “share-alike” requirement and BSD’s permission to go proprietary. Her models feature finitely many generations of developers, rendering open source licenses in her models not self-referential.⁶

In comparison to this strand of studies, our model has the following two properties. First, it is rich enough that the defining features of GPL (its “share-alike” requirement) and BSD (its permission to go proprietary) can be described explicitly, and hence the optimality of mixing-and-matching these features can be meaningfully discussed (in contrast to Lerner and Tirole, 2005b). Second, it is a model with infinitely many generations of developers, and hence open source licenses in the model are self-referential (in contrast to Gaudeul, 2004, 2005).

⁶For example, the penultimate generation’s licenses do not restrict the ultimate generation’s choices of licenses.

2 The Basic Model

In this section, we shall first describe a basic model that is barely rich enough to accommodate GPL and BSD. We will explain how this basic model can be extended to accommodate other open source licenses in subsequent sections.

Consider a discrete-time model, where time is indexed by $t = 0, 1, 2, \dots$. In every period t , there is one and only one developer, called developer t , who has the potential of developing a software, called software t . Note that we abuse notation by using the same index, t , for time, for developer, and for software.

We can think of developer 0 as an original developer such as Linus Torvalds or Donald Knuth, and software 0 his original software. For any $t > 0$, we can think of software t as a more advanced version of software $t - 1$, perhaps by adding extra functionalities. Of course, indirectly via software $t - 1$, software t is also a more advanced version of any software $s < t - 1$.

Developer t will be able to develop software t only if he has an opportunity to learn from the source code of software $t - 1$. Apparently, this cannot happen if developer $t - 1$ chose to go proprietary (i.e., to sell copies of the binary code for a profit), instead of sharing the source code with the others. Therefore, we shall assume that, if developer $t - 1$ goes proprietary, none of the software $s \geq t$ can be developed. In other words, the game ends after developer $t - 1$ goes proprietary.

We assume that if developer $t - 1$ did not go proprietary, then he must go open source; i.e, sharing his source code with the others using one of the open source licenses. In other words, keeping the software private is not an option. This is not a strong assumption. When a developer is not going to make a profit out of his creation, any tiny preference of sharing will prompt him to share.

If developer $t - 1$ chose to go open source, thus enabling developer t to develop software t , the game continues to period t . In period t , developer t first draws a tuple of four variables, $\theta_t = (c_t, \pi_t, w_t, W_t)$, where c_t is his cost of developing software t , π_t and w_t are the potential profit and consumer surplus, respectively, if he develops software t and then goes proprietary, and W_t is the consumer surplus

if he develops software t and then goes open source instead.⁷ The tuple θ_t is drawn from the probability law $P(\theta_t | \theta^{t-1})$, where $\theta^{t-1} = (\theta_0, \dots, \theta_{t-1})$. We assume that the realization of θ_t is observable to all developers $s \geq t$.

Upon observing θ_t , developer t then decides whether to pay the development cost c_t and develop software t . If he decides not to (an option denoted by **Q**, which stands for quitting), the game ends.

If he decides to develop software t , he then decides whether to go proprietary or to go open source using one of the open source licenses. Some of these options may not be available, depending on the open source license chosen by developer $t - 1$. We first enumerate the three different options developer t may have, and then explain which subsets of these options are available to him given different open source licenses chosen by developer $t - 1$.

P: *going proprietary*

Developer t goes proprietary, realizing potential profit π_t and consumer surplus w_t , and the game ends.

G: *going open source using the GPL license*

Developer t goes open source, thus enabling developer $t + 1$ to develop software $t + 1$. However, if developer $t + 1$ chooses to develop software $t + 1$, he has to go open source using GPL (i.e., choosing **G**) as well. For developer t , potential profit π_t is forgone, and the consumer surplus is W_t .

B: *going open source using the BSD license*

Developer t goes open source, thus enabling developer $t + 1$ to develop software $t + 1$. If developer $t + 1$ chooses to develop software $t + 1$, he can either go proprietary (i.e., choosing **P**), or go open source using any of the two licenses (i.e., choosing **G** or **B**). For developer t , potential profit π_t is forgone, and the consumer surplus is W_t .

⁷It is natural to assume that $W_t > w_t$. None of our results, however, require this assumption. All of our examples, meanwhile, satisfy this extra assumption.

We should hasten to emphasize that what we called the consumer surplus (either w_t or W_t) should more appropriately be understood as the *internalizable* part of the consumer surplus. It is the part that current and past developers (i.e., developers $s \leq t$) who went open source can internalize. It is likely only a small part of the whole consumer surplus. In particular, we do not presume that $W_t > \pi_t + w_t$, which would have been a natural inequality to assume had we interpreted these as the whole consumer surplus that can potentially be generated. Indeed, allowing for $W_t < \pi_t + w_t$ is necessary to explain why developer t may sometimes go proprietary.⁸

For developer $t > 0$, which of these three options (**P**, **G**, and **B**) are available depends on what open source license he is subject to (i.e., what open source license developer $t - 1$ chose). We enumerate these different situations in Table 1.

what developer $t - 1$ chose	what developer t can choose after developing software t
P	the game ended in period $t - 1$ already
G	G
B	P, G, and B

Table 1: The basic model.

For developer 0, we assume that all three options (**P**, **G**, and **B**) are available to him, and hence we may treat him as if he is subject to license **B**.

Let T be the period when the game ends. The game ends in period T iff (1) $\forall t < T$, developer t decided to develop software t and then went open source, and (2) developer T either (2a) decides not to develop software T , or (2b) decides to develop software T and then goes proprietary. If the game never ends, let $T = \infty$.

For most of this paper (except for Subsection 7.2), we assume exponential

⁸While our favorite interpretation of w_t and W_t is the (internalizable part of the) consumer surplus, this is not the only possible interpretation. Alternatively, one can interpret them as what Raymond (1999) calls “egoboo”; i.e., the ego boost all developers $s \leq t$ can gain from their increased visibility and recognition when the more advanced software t is developed.

discounting. For any developer $t < T$, his payoff is

$$U_t := -c_t + \sum_{s=t}^{T-1} \beta^{s-t} W_s + \beta^{T-t} \times \begin{cases} 0 & \text{if developer } T \text{ does not develop software } T \\ w_T & \text{if developer } T \text{ develops software } T \text{ and goes proprietary} \end{cases} \quad (1)$$

where $\beta \in (0, 1)$ is the common discount factor. In other words, we assume that if developer t ever incurs the development cost c_t but then forgoes his potential profit π_t , he does so because he is altruistic enough to care about current and future consumer surplus.⁹

As for developer T , his utility is 0 if he decides not to develop software T , or is $\pi_T + w_T - c_T$ if he decides to develop software T and then goes proprietary.

We assume that developer t 's choice is observable to all developers $s > t$. We have thus described an infinite-horizon observable-action game. The primitives of the game are the probability law $P(\cdot | \cdot)$ and the common discount factor β . Our solution concept is subgame perfect equilibrium.

It is easy to see that, depending on the probability law $P(\cdot | \cdot)$, it can be strictly optimal for developer 0 to choose each of the four options (**Q**, **P**, **G**, and **B**). For example, he would **Quit** if c_0 is prohibitively large; goes **Proprietary** if W_0 is small compared to $\pi_0 + w_0$ and all (w_t, W_t) , $t \geq 1$, are likely to be negligible; goes open source using **GPL** if $\pi_0 + w_0$ is small compared to W_0 and all w_t , $t \geq 1$, are likely to be negligible; and goes open source using **BSD** if $\pi_0 + w_0$ is small compared to W_0 , w_1 is likely to be similar to W_1 , and all W_t , $t \geq 2$ are likely to be negligible.

Without imposing significantly more structure on the probability law $P(\cdot | \cdot)$, it seems difficult to fully characterize the conditions under which developer 0 would choose each of these options. Since our main focus is on what *other* open source licenses are available, we shall not pursue such a full characterization here, and

⁹One may wonder whether the consumer surplus W_t generated by software t depends on whether software $t + 1$ will be developed, as the latter may supersede the former. In that case we may reinterpret W_{t+1} as the marginal increase in total consumer surplus generated by the creation of software $t + 1$, taking into account the creative destruction effect on software t (as well as all the earlier softwares).

shall leave it for future research. Appendix A, however, demonstrates how, if we are willing to impose strong assumptions on the probability law $P(\cdot|\cdot)$, some interesting characterizations can be obtained.

3 A Simple Extension

The basic model in Section 2 is so flexible that adding new open source licenses is easy. For example, it is easy to see how anyone, by playing with Table 1, would naturally come up with the new licenses depicted in Table 2.¹⁰ In Table 2, **R** is simply the mirror image of **G**. It refers to itself in its own definition, in exactly the same self-referential manner as **G**. But it deviates from **G** by allowing the next developer to go proprietary. License **1** is derived from **G** in a different way. It puts the same restriction on the next developer's choices of open source licenses, but deviates from **G** by also allowing the next developer to go proprietary. Finally, license **L** is derived from **R** in exactly the same way as how **1** is derived from **G**. It puts the same restriction on the next developer's choices of open source licenses, but deviates from **R** by *not* allowing the next developer to go proprietary.

what developer $t - 1$ chose	what developer t can choose after developing software t
G	G
R	P and R
1	P and G
L	R

Table 2: A simple extension.

We call **R** the *recursive-BSD license*. It pushes the defining feature of BSD (its permission to go proprietary) to its limit. Not only that the next developer is allowed to go proprietary, all future developers, to the extent that they have a chance to move, are allowed to go proprietary as well. This is because, for example,

¹⁰These authors confess that they had lots of fun filling out Table 2.

the next developer cannot use GPL to forbid the next-next developer from going proprietary. In this sense, **R** is even more permissive than BSD. Of course, one can also argue that **R** is more restrictive than BSD, as it puts more restrictions on how the next developer can restrict the next-next developer. This is a perfect example of why it can be difficult to summarize an open source license by a single parameter called “permissiveness” as in Lerner and Tirole (2005b).

We call license **1** the *1-chance-only license*. License **1** is unambiguously more restrictive than **R**. While it, like **R**, gives the next developer the permission to go proprietary, it forbids the next developer from giving the same permission to the next-next developer. In other words, license **1** gives future developers one and only one chance to go proprietary, namely in the next period and in the next period only.

We call license **I** the *1-time-forbiddance license*. As suggested by the symbol, license **I** in effect turns the idea of license **1** upside-down. While license **1** allows the next developer to go proprietary but forbids him and any future developer from allowing their successors to do the same, license **I** does not allow the next developer to go proprietary but requires that he and any future developer allow their successors to do so.

Are these new licenses bogus? Are there any reasons why an open source developer may consider using them? Are there situations where they may serve purposes that neither GPL nor BSD serve? The answer is a qualified “no”. Indeed, in Section 6, we shall prove that, within our current model setup, these new licenses can all be ignored without loss of generality (see Theorem 2). This is a surprising result, given that the proof requires almost no restrictions on the probability law $P(\cdot|\cdot)$ except for absolute continuity. This is the strongest explanation to date of why GPL and BSD stood out from other licenses as the most popular two among open source developers.

However, that our “no” answer needs to be qualified also means that it is not universally true. In Section 7, we shall demonstrate how these new licenses can

beat GPL and BSD once we deviate from our current model setup. Specifically, in Subsection 7.1, we show how license **R** can beat GPL and BSD if developers form a tree instead of a linear structure. In Subsection 7.2, we show how license **1** can beat GPL and BSD if developers have hyperbolic instead of exponential discounting.

Before we can formally state our sufficiency and insufficiency results in Sections 6 and 7, respectively, we have to first introduce two more concepts: irreducibility (Section 4), and imposture proofness (Section 5).

4 Irreducible Space of Open Source Licenses

The new licenses introduced in Section 3 are not the only possible new open source licenses. Indeed, a little thought would suggest that infinitely many can be generated in a similar manner. It is hence important to have a sense of how the set of all possible open source licenses looks like, and then impose some intuitive axioms to weed out the less interesting ones.

We can define a general space of open source licenses in a manner similar to that in Sections 2 and 3. For any set X , let $\mathcal{P}(X)$ be the set of all *nonempty* subsets of X .¹¹

Definition 1 *Let O be an arbitrary set, with each element $o \in O$ corresponding to an open source license. Let $g : O \rightarrow \mathcal{P}(\{\mathbf{P}\} \cup O)$ be a nonempty correspondence such that, for any $o \in O$, $g(o) \cap O \neq \emptyset$. Then, $S = (O, g)$ is a space of open source licenses.*

Intuitively, $g(o)$ specifies the (nonempty) set of options available to a developer who is subject to open source license o .¹² We assume that going open source (using *some* open source license) must always be an option, which translates

¹¹In mathematics, the notation $\mathcal{P}(X)$ typically stands for the power set of X , which is the set of all subsets of X , including the empty set. Here, it will ease our notation if we re-define $\mathcal{P}(X)$ as the set of all *nonempty* subsets of X instead.

¹²It goes without saying that quitting, **Q**, is always an option. We hence omit it from the developer's choice set for brevity.

into the requirement that $g(o) \cap O \neq \emptyset$. This assumption is natural, as it seems impossible to design any open source license that forces the next developer to go proprietary.¹³

A space of open source licenses may contain duplicates of otherwise identical open source licenses. For example, consider $O = \{G_1, G_2\}$, with $g(G_1) = \{G_2\}$ and $g(G_2) = \{G_2\}$. Then the differences between G_1 and G_2 are superfluous, and one may for all purposes regard both as being identical to GPL.

Formally, for any space of open source license $S = (O, g)$, let \sim be an equivalence relation on O .¹⁴ For any $o \in O$, let $[o]$ denote the equivalence class of o . We shall abuse notation by sometimes writing P as $[P]$ as well. Let O^\sim denote the corresponding quotient set (i.e., the set of equivalence classes of elements in O). Let $\mu : O \cup \{P\} \rightarrow O^\sim \cup \{P\}$ be the canonical mapping such that, $\forall x \in O \cup \{P\}$, $\mu(x) = [x]$. For any $o \in O$, let $(\mu \circ g)(o) = \{[x] \in O^\sim \cup \{P\} : x \in g(o)\}$, which is nonempty because $g(o)$ is nonempty. We say that g and \sim are *compatible* if $(\mu \circ g)(o) = (\mu \circ g)(o')$ whenever $o \sim o'$. If g is compatible with \sim , we can define $g^\sim : O^\sim \rightarrow \mathcal{P}(O^\sim \cup \{P\})$ such that $g^\sim([o]) = (\mu \circ g)(o)$ for any $o \in O$. Then $S^\sim = (O^\sim, g^\sim)$ is a space of open source licenses. We say that \sim is nontrivial if there exist distinct $o, o' \in O$ such that $o \sim o'$. If there exists a nontrivial equivalence relation \sim compatible with g , we say that $S = (O, g)$ is *reducible* to $S^\sim = (O^\sim, g^\sim)$ (or simply *reducible*); otherwise $S = (O, g)$ is *irreducible*.

Consider our earlier example where $O = \{G_1, G_2\}$, $g(G_1) = \{G_2\}$, and $g(G_2) = \{G_2\}$. The only nontrivial equivalence relation is such that $G_1 \sim G_2$. This equivalence relation is compatible with g , because $(\mu \circ g)(G_1) = \{[G_2]\} = (\mu \circ g)(G_2)$. Therefore, $S = (O, g)$ is reducible to $S^\sim = (O^\sim, g^\sim)$, where $O^\sim = \{[G_2]\}$ and $g^\sim([G_2]) = \{[G_2]\}$.

¹³It seems impossible to stop the next developer from giving up his copyright and putting his source code in the public domain, which in our model is equivalent to going open source using BSD. It should however be pointed out that, in reality, there are subtle differences between “putting the source code in the public domain” and “going open source with the BSD licenses”, and these subtle differences are not captured by our model.

¹⁴A binary relation \sim is an equivalence relation if it is reflexive, symmetric, and transitive.

Consider the basic model in Section 2 where $O = \{\mathbf{G}, \mathbf{B}\}$, $g(\mathbf{G}) = \{\mathbf{G}\}$, and $g(\mathbf{B}) = \{\mathbf{P}, \mathbf{G}, \mathbf{B}\}$. The only nontrivial equivalence relation is such that $\mathbf{G} \sim \mathbf{B}$. This equivalence relation is not compatible with g , however, because $(\mu \circ g)(\mathbf{G}) = \{[\mathbf{G}]\} \neq \{\mathbf{P}, [\mathbf{G}]\} = (\mu \circ g)(\mathbf{B})$. Therefore, $\mathcal{S} = (O, g)$ is irreducible.

One interesting question is whether there exists a universal space of open source licenses, $\mathcal{S}^U = (O^U, g^U)$, that contains every space of open source licenses as its sub-space. The answer is yes. The construction of such a universal space is similar to that of the universal type space in epistemic game theory. We include this construction in Appendix B for completeness.

5 Imposture-Proof Open Source Licenses

To motivate the idea of imposture-proofness, let's revisit the 1-time-forbiddance license introduced in Section 3:

L *going open source with the 1-time-forbiddance license*

Developer t goes open source. If developer $t + 1$ chooses to develop software $t + 1$, he has to go open source with the recursive-BSD license (i.e., choosing **R**), thus enabling future developers to go proprietary.

Imagine that developer $t - 1$ goes open source with license **L**, hoping to prohibit developer t from going proprietary. One possible way for developer t to game this license is to split his software t into two successive versions, version $t.1$ and version $t.2$, with version $t.1$ a more advanced version of software $t - 1$, and version $t.2$ a more advanced version of version $t.1$. He can roll out version $t.1$ first, go open source with license **R**, thus satisfying the terms in developer $t - 1$'s license **L**. He can then roll out version $t.2$, possibly using a different identity, and then go proprietary, which is allowed by the terms in version $t.1$'s license **R**.

The reason why developer t can game license **L** is that, while **L** precludes **P**, it allows for **R** which in turn allows for **P**. More generally, any license that tries to

preclude option x but allows for some open source license o that allows for option x can be gamed in a similar manner. This motivates the following axiom.

Definition 2 Let $\mathcal{S} = (\mathcal{O}, g)$ be a space of open source licenses. An open source license $o \in \mathcal{O}$ is said to be *imposture-proof* if

$$o' \in g(o) \cap \mathcal{O} \implies g(o') \subseteq g(o).$$

The space \mathcal{S} is said to be *imposture-proof* if every open source license $o \in \mathcal{O}$ is imposture-proof.

Note that all the other open source licenses studied in Section 3 (i.e., **G**, **R**, and **1**) are imposture-proof open source licenses.

The following theorem is the main result of this section that we shall use in Section 6.

Theorem 1 Let $\mathcal{S} = (\mathcal{O}, g)$ be an irreducible space of open source licenses that is imposture-proof. Any open source license $o \in \mathcal{O}$ that does not allow the next developer to go proprietary (i.e., $\mathbf{P} \notin g(o)$) is identical to the GPL license **G** in the sense that $g(o) = \{o\}$.

PROOF: Consider any open source license $o \in \mathcal{O}$ such that $\mathbf{P} \notin g(o)$. If $o' \notin g(o)$ for any $o' \neq o$, then by the nonemptiness of $g(o) \cap \mathcal{O}$ we must have $g(o) = \{o\}$, and we are done. Therefore, let's suppose there exists $o' \neq o$ such that $o' \in g(o)$. Let's define an equivalence relation \sim such that $[o] = g(o) \cup \{o\}$, and $[o''] = \{o''\}$ for any $o'' \notin g(o) \cup \{o\}$. This is a nontrivial equivalence relation because $o' \neq o$ and yet $o' \sim o$. We shall prove that g and \sim are compatible, and hence $\mathcal{S} = (\mathcal{O}, g)$ is reducible. To prove compatibility, it suffices to prove that $(\mu \circ g)(o') = (\mu \circ g)(o)$ for any $o' \in g(o)$. By imposture-proofness, we have

$$g(o') \subseteq g(o) \subseteq g(o) \cup \{o\} = [o],$$

and hence $(\mu \circ g)(o') = \{[o]\} = (\mu \circ g)(o)$ as claimed. \blacksquare

In Appendix C, we shall categorize different imposture-proof open source licenses. This categorization is independent of the results in Section 6, which deals with the question of why GPL and BSD stood out from other open source licenses as the most popular two among open source developers. The categorization, however, may be useful for future research, especially in light of our results in Section 7, which demonstrate how open source licenses other than GPL and BSD can be useful if we go beyond our current model setup.

6 Sufficiency of GPL and BSD

In this section, we shall prove that, in any subgame perfect equilibrium, if developer 0 is ever going to go open source, he cannot do better than going open source with either GPL or BSD. This result is of interest because it sheds light on why GPL and BSD stood out from other open source licenses as the most popular two among open source developers.

Theorem 2 *Assume that the probability law $P(\cdot|\cdot)$ is absolutely continuous. Let $\mathcal{S} = (\mathcal{O}, g)$ be an irreducible space of open source licenses that is imposture-proof and contains open source licenses $\mathbf{G}, \mathbf{B} \in \mathcal{O}$ such that $g(\mathbf{G}) = \{\mathbf{G}\}$ and $g(\mathbf{B}) = \{\mathbf{P}\} \cup \mathcal{O}$. Then, in any subgame perfect equilibrium, if developer 0 is ever going to go open source, he cannot do better than going open source using either \mathbf{G} or \mathbf{B} .*

PROOF: We have already argued in Section 2 that either \mathbf{G} or \mathbf{B} can be strictly optimal for developer 0 if these are the only two open source licenses available to him. Since $\mathcal{S} = (\mathcal{O}, g)$ is irreducible and imposture-proof, by Theorem 1, any open source license $o \in \mathcal{O}$ other than \mathbf{G} must have the property that $\mathbf{P} \in g(o)$. Therefore, it suffices to prove that, in any subgame perfect equilibrium, choosing \mathbf{B} is weakly better than choosing any of these open source licenses for developer 0.

Consider any open source license $o \in \mathcal{O}$ such that $\mathbf{P} \in g(o)$. Suppose developer 0 goes open source using license o . For any $o_1 \in g(o) \cap \mathcal{O}$, let $v_1(o_1)$ be developer 1's gross expected payoff (gross of development cost c_1) if he goes open source with o_1 , where the expectation is taken over the realizations of $\{\theta_t\}_{t \geq 2}$, and is taken conditional on the equilibrium strategies of developers $t \geq 2$. Note that $v_1(o_1) = W_1 + \beta \times (\dots) \geq W_1$. Let $v_1^* = \sup_{o_1 \in g(o) \cap \mathcal{O}} v_1(o_1)$. Developer 1's optimal strategy is hence

1. not to develop software 1 if $\theta_1 \in E_1^0 := \{\theta_1 \mid c_1 > \max\{\pi_1 + w_1, v_1^*\}\}$;
2. to develop software 1 and go proprietary if $\theta_1 \in E_1^P := \{\theta_1 \mid \pi_1 + w_1 > \max\{c_1, v_1^*\}\}$; and
3. to develop software 1 and go open source with one of the open source licenses in $\operatorname{argmax}_{o_1 \in g(o) \cap \mathcal{O}} v_1(o_1)$ if $\theta_1 \in E_1^* := \{\theta_1 \mid v_1^* > \max\{c_1, \pi_1 + w_1\}\}$.¹⁵

Since $P(\cdot \mid \cdot)$ is absolutely continuous, the boundaries of events E_1^0 , E_1^P , and E_1^* have probability 0 and hence can be ignored.

By going open source with o , developer 0's gross expected payoff (gross of development cost c_0) is hence

$$W_0 + \beta \mathbb{E} \left[0 \times \mathbf{1}_{E_1^0} + w_1 \times \mathbf{1}_{E_1^P} + v_1^* \times \mathbf{1}_{E_1^*} \mid \theta_0 \right], \quad (2)$$

where $\mathbf{1}_E$ is the indicator function of event E , and $\mathbb{E}[\cdot]$ is the expectation operator.

Note that v_1^* is weakly increasing in the set $g(o) \cap \mathcal{O}$ (according to the order of set inclusion), and is maximized at \mathbf{B} . When v_1^* increases, there are three effects on the expression (2):

1. an increase in v_1^* directly increases the expression (2);
2. some θ_1 move from event E_1^0 to event E_1^* , which weakly increases the expression (2) because any such θ_1 must feature $v_1^* \geq 0$;

¹⁵The sup v_1^* is attainable in this last case because an optimal strategy exists in a subgame perfect equilibrium.

3. some θ_1 move from event E_1^P to event E_1^* , which weakly increases the expression (2) because any such θ_1 must feature $v_1^* \geq \pi_1 + w_1 \geq w_1$.

This proves that, in any subgame perfect equilibrium, choosing **B** is weakly better than choosing any open source license $o \in \mathcal{O}$ such that $\mathbf{P} \in g(o)$. ■

In the proof of Theorem 2, the assumption of an absolutely continuous probability law $P(\cdot | \cdot)$ implies that developer 1 is indifferent (between developing software 1 or not, and between going proprietary or going open source) with probability 0. If developer 1 is indifferent with strictly positive probability, how he breaks ties has non-trivial implications on developer 0's gross expected payoff. If developer 1 breaks ties in a manner that depends on some payoff-irrelevant details—such as the name of the license, or whether the license allows for certain irrelevant options—then it is possible that BSD does not fare as well as another license simply because the former leads developer 1 to break ties in a way that is unfavorable to developer 0. If we alternatively assume that, say, developers always break ties in favor of developing the software, and in favor of going open source, then we can relax the assumption of absolute continuous $P(\cdot | \cdot)$ in Theorem 2.

Theorem 2 is silent on the existence of an equilibrium. It is easy to construct simple examples where equilibria exist, and hence Theorem 2 is definitely not a characterization of the empty set. It is also easy to construct (not necessarily interesting) sufficient conditions for existence of an equilibrium. But the applicability of Theorem 2 goes beyond the class of games confined by these sufficient conditions. It says that whenever an equilibrium exists, regardless of whether these sufficient conditions are satisfied, developer 0 cannot do worse by ignoring open source licenses other than GPL and BSD.

7 Insufficiency of GPL and BSD

The sufficiency result in Section 6 is surprising, especially given how little restriction on the probability law $P(\cdot|\cdot)$ it relies on. Nevertheless, the proof of Theorem 2 does rely on two particular assumptions that we have maintained up to this point. The first is a linear structure of developers. The second is exponential discounting. These two assumptions together imply that the interests of developers 0 and 1 are aligned conditional on the event that developer 1 goes open source. This alignment explains why developer 0 would like to give developer 1 the maximum freedom to choose which open source license to license software 1.

In this section, we shall show how the sufficiency result in Section 6 breaks down once we relax any one of these two assumptions. Specifically, in Subsection 7.1, we shall show how license **R** can beat GPL and BSD if developers form a tree instead of a linear structure. In Subsection 7.2, we shall show how license **1** can beat GPL and BSD if developers have hyperbolic instead of exponential discounting.

7.1 An Example with a Tree Structure

Suppose developer 0, if going open source, can inspire two instead of only one developers — call them developers 1A and 1B, respectively. Suppose this is the only time that an open source software can inspire two developers. It means that we will have a linear structure again starting from developer 1A: $\forall n = 1, 2, \dots$, an open source software nA inspires one and only one developer, namely developer $(n + 1)A$, and so on. Likewise, there is a linear structure starting from developer 1B.

It is natural to assume that every developer cares only about the consumer surplus generated by the offsprings of his own software. Therefore, for example, developer 1A cares only about the consumer surplus generated by softwares 1A, 2A, 3A. . . , and developer 1B cares only about the consumer surplus generated by

softwares 1B, 2B, 3B. . . , while developer 0 cares about consumer surplus generated by all softwares. This creates a new wedge between the interests of developers 0 and 1A, for example, even conditional on the event that developer 1A develops software 1A and goes open source. This wedge results in a breakdown of the proof of Theorem 2.

Suppose $S = \{O, g\}$ is such that $O = \{G, R, B\}$, where

$$\begin{aligned} g(G) &= \{G\}, \\ g(R) &= \{P, R\}, \\ \text{and } g(B) &= \{P\} \cup O. \end{aligned}$$

Suppose developer 0 draws $(c_0, \pi_0, w_0, W_0) = (0, 0, 0, 1)$ with certainty. Apparently, he strictly prefers to develop software 0 and then go open source. Knowing that his open source software will for sure inspire two developers, but not knowing their identities *ex ante*, he uses a single open source license $o \in O$ to license his software to both of them.

Suppose softwares 1A and 1B are “independent” of each other; roughly speaking, it means they neither complement nor substitute each other. Specifically, each of developers 1A and 1B draws $(c_1, \pi_1, w_1, W_1) = (0, 0, 0, 1)$ with certainty, where $W_1 = 1$, for example, is the (internalizable) (marginal increase in) consumer surplus generated by an open source software 1A (1B) regardless of whether software 1B (1A) is developed and how it is distributed. Apparently, developer 1A (1B) strictly prefers to develop software 1A (1B) and then go open source. What open source licenses he can choose, however, depends on developer 0’s license o .

We assume that developers 1A and 1B choose their licenses o_A and o_B , respectively, simultaneously. Their choices are then observable to both developers 2A and 2B.

Most of the actions in this example will take place in the game between developers 2A and 2B. Specifically, we “truncate” the game after generation 2 as follows:

$\forall n \geq 3$, developer nA (nB) draws $(c_n, \pi_n, w_n, W_n) = (1, 0, 0, 0)$ with certainty, and hence no developers after generation 2 will bother to develop their respective softwares even if they have the chance to do so.

Before we describe the game between developers 2A and 2B in details, let's first sketch the main idea behind the construction. Softwares 2A and 2B are imperfect substitutes. The two softwares together can generate more consumer surplus than, but not as much as two times of, what a single software can generate. It means the consumer surplus attributable to software 2B, for example, is smaller in the presence of software 2A, and gets even smaller if 2A goes open source. When developer 1A forces software 2A to go open source, say by using license **G**, he generates a negative externality on developer 1B by reducing the consumer surplus attributable to software 2B. While developer 1A does not internalize this negative externality, developer 0 does. Developer 0 hence may want to choose a license that forbids developer 1A from using license **G** and generating this negative externality. License **R** serves exactly this purpose.

We now describe the game between developers 2A and 2B. The game is symmetric, so we describe only the half from the perspective of developer 2A. Specifically, developer 2A's (W, w, π) depends on the presence and the mode of distribution of software 2B as in Table 3.

		conditional on 2B		
		quitting	going proprietary	going open source
2A's	W	100	90	60
	w	50	40	20
	π	200	160	70

Table 3: 2A's (W, w, π) conditional on 2B's action.

The first notable feature of Table 3 is that, as explained earlier, the consumer surplus attributable to software 2A is smaller in the presence of software 2B, and gets even smaller if 2B goes open source. The second notable feature is that developer 2A always strictly prefers going proprietary to going open source (i.e.,

$w + \pi > W$) regardless of what developer 2B does (recall that their generation is effectively the last generation). Only a license **G** from developer 1A will be able to force developer 2A to go open source. Since the game is symmetric, Table 3 is also the table of developer 2B's (W, w, π) conditional on the presence and the mode of distribution of software 2A.

The developing costs, c_A and c_B , of developers 2A and 2B, respectively, are drawn from the joint distribution depicted in Table 4.

		c_B	
		85	0
c_A	85	0.4	0.3
	0	0.3	0

Table 4: Joint distribution of c_A and c_B .

Note that, when $c_A = 0$, developer 2A strictly prefers to develop software 2A; and likewise for the case of $c_B = 0$. When $c_A = 85$, he will still strictly prefers to develop if he is allowed (say by a licence **R**) to go proprietary, because $w + \pi$ is always strictly bigger than 85 regardless of what developer 2B does. It remains to study his entry decision if he is only allowed to go open source (by a license **G**). We consider two cases.

The first case is that his opponent, developer 2B, is allowed to go proprietary (say by a license **R**). In this case, as we have argued earlier, developer 2B will for sure develop and go proprietary regardless of c_B . Therefore, by developing and going open source, developer 2A can obtain $W = 90 > 85$ according to Table 3. He will hence develop even when $c_A = 85$.

The second case is that his opponent, developer 2B, is also only allowed to go open source (by a license **G**). In this case, as we have argued earlier, developer 2B will for sure develop and go open source if $c_B = 85$. This in itself is bad enough news for developer 2A, because it happens with conditional probability $3/7$ when $c_A = 85$. Even in the best case scenario where developer 2B quits when $c_B = 0$, developer 2A's conditional expected W is still only $(4/7) \times 100 + (3/7) \times 60 < 85$

when $c_A = 85$. Therefore, he will not develop when $c_A = 85$.

We summarize the entry decisions of developers 2A and 2B, conditional on the licenses they are subject to, in Table 5. We do not include license **B** in the table, because it is effectively the same as license **R** from the perspective of these developers (recall that they are effectively the last generation).

		if 2B is subject to license	
		R	G
if 2A is subject to license	R	yes \ yes	yes \ yes
	G	yes \ yes	no \ no

Table 5: Will 2A (2B) develop if $c_A = 85$ ($c_B = 85$)?

We now use Table 5 to study the licensing game between developers 1A and 1B. The game will be trivial if developer 0 used license **G**, in which case developers 1A and 1B have no choice but to use license **G** as well. According to Table 5, developer 2A (2B) will develop software 2A (2B) and go open source iff $c_A = 0$ ($c_B = 0$). According to Table 4, it will never happen that both softwares 2A and 2B are developed. Each software will be developed with probability 0.3. And whenever it is developed, it is developed alone, generating consumer surplus 100. The expected consumer surplus attributable to software 2A (2B) and hence internalized by developer 1A (1B) is hence $0.3 \times 100 = 30$, whereas developer 0 internalizes the sum 60.

The game will also be trivial if developer 0 used license **R**, in which case developers 1A and 1B have no choice but to use license **R** as well (recall that they both strictly prefer to go open source rather than proprietary). According to Table 5, developer 2A (2B) will always develop software 2A (2B) and go proprietary regardless of c_A (c_B). The market will hence always have two proprietary generation-2 softwares available. According to Table 3, each proprietary software generates consumer surplus 40, which is internalized by developers 1A and 1B, whereas developer 0 internalizes the sum 80.

The game is non-trivial only if developer 0 used license **B**, in which case both

developers 1A and 1B have the freedom to choose between licenses **R** and **G**.¹⁶ We shall argue that using license **G** is the dominant strategy for both developers.

First consider the case where opponent 1B uses license **R**. In this case, according to Table 5, whether 1A uses license **R** or **G** will not affect developer 2A's entry decision. Since, according to Table 3, $W > w$ regardless of what developer 2B does, developer 1A strictly prefers that developer 2A goes open source, and hence strictly prefers to respond with license **G** rather than **R**.

Then consider the case where opponent 1B uses license **G**. If 1A uses license **R**, according to Table 5, both softwares 2A and 2B will be developed for sure, with 2A going proprietary and 2B going open source. According to Table 3, the consumer surplus generated by the proprietary software 2A will be 20. If 1A uses license **G** instead, according to the analysis two paragraphs earlier, the expected consumer surplus generated by software 2A will be 30. Therefore, once again, developer 1A strictly prefers to respond with license **G** rather than **R**.

To summarize, if developer 0 used license **B**, the unique continuation equilibrium is such that both developers 1A and 1B play their dominant strategy of using license **G**. This is the same outcome as in the case had developer 0 used license **G**. In both cases the total expected consumer surplus generated by generation-2 softwares is 60. If developer 0 uses license **R** instead, the total consumer surplus generated by generation-2 softwares will be 80. Developer 0 hence strictly prefers license **R** to licenses **B** and **G**.

Proposition 1 *In the example in this subsection, there is a unique subgame perfect equilibrium, where developer 0 develops software 0 and goes open source with license **R**, developers 1A and 1B both develop their respective softwares and go open source with license **R**, and developers 2A and 2B both develop their respective softwares and go proprietary. Specifically, developer 0 strictly prefers using license **R** rather than licenses **B** and **G**, even though the latter are also available.*

¹⁶Again, we ignore the choice of license **B**, which is effectively the same as license **R** from the perspective of developers 2A and 2B, who are effectively the last generation.

7.2 An Example with Hyperbolic Discounting

It has become a common practice in behavioral economics to *approximate* hyperbolic discounting using an (α, β) -formulation. Here, we shall instead adopt an $(\alpha_1, \alpha_2, \beta)$ -formulation. Specifically, given any infinite sequence of dated payoffs, $\{u_t, u_{t+1}, u_{t+2}, \dots\}$, we assume that developer t 's present discounted value is

$$U_t = u_t + \alpha_1 \left(u_{t+1} + \alpha_2 \left(u_{t+2} + \beta u_{t+3} + \beta^2 u_{t+4} + \dots \right) \right),$$

where $\alpha_1 \leq \alpha_2 \leq \beta$.¹⁷

Consider the following example with hyperbolic discounting. Suppose $\alpha_1 = 1/4$, $\alpha_2 = 2/4$, and $\beta = 3/4$; and $\forall t \geq 0$, $w_t \equiv w = 15$, and $W_t \equiv W = 20$. Suppose the joint distribution of c_t and π_t is iid across t , and is as depicted in Table 6.

$P(c_t, \pi_t)$	$\pi_t = 0$	$\pi_t = 20$	$\pi_t = 35$
$c_t = 10$	$\epsilon \approx 0$	0	0
$c_t = 20$	0	$(1 - \epsilon)/2$	0
$c_t = 40$	0	0	$(1 - \epsilon)/2$

Table 6: Joint distribution of c_t and π_t .

Suppose $S = \{O, g\}$ is such that $O = \{\mathbf{G}, \mathbf{R}, \mathbf{1}, \mathbf{B}\}$, where

$$\begin{aligned} g(\mathbf{G}) &= \{\mathbf{G}\}, \\ g(\mathbf{R}) &= \{\mathbf{P}, \mathbf{R}\}, \\ g(\mathbf{1}) &= \{\mathbf{P}, \mathbf{G}\}, \\ \text{and } g(\mathbf{B}) &= \{\mathbf{P}\} \cup O. \end{aligned}$$

Since this is a stationary environment, we shall further restrict our attention to pure-strategy Markov perfect equilibrium, in which every developer t follows the

¹⁷As will be explained later, α_1 actually does not play any role, and hence the restriction $\alpha_1 \leq \alpha_2$ is redundant. We maintain this restriction only to stay in sync with the hyperbolic-discounting literature.

same pure Markov strategy that depends only on (i) the open source license he is subject to, which in turn is chosen by developer $t - 1$,^{18,19} and (ii) the realization of (c_t, π_t) . More formally, a pure Markov strategy is a function $\sigma : (c, \pi, o) \mapsto \{\mathbf{Q}, \mathbf{P}\} \cup \mathcal{O}$ that describes any developer's choice over $\{\mathbf{Q}, \mathbf{P}\} \cup \mathcal{O}$ given his draw of (c, π) and the open source license $o \in \mathcal{O}$ that he is subject to, with the restriction that $\forall (c, \pi, o)$, $\sigma(c, \pi, o) \in \{\mathbf{Q}\} \cup g(o)$. A pure-strategy Markov perfect equilibrium (hereafter, an *equilibrium*) is a Markov strategy σ^* that is optimal for any developer t among all pure Markov strategies if all future developers $s > t$ follow the pure Markov strategy σ^* .

Consider any developer t . If he develops software t and goes open source, his gross payoff (gross of development cost c_t) is at least $W_t = W = 20$ (which is achieved if no future software is developed), and is at most

$$W_t + \alpha_1 \left(W_{t+1} + \alpha_2 \left(W_{t+2} + \beta W_{t+3} + \beta^2 W_{t+4} + \dots \right) \right) = W + \alpha_1 \left(W + \alpha_2 \left(W + \beta W + \beta^2 W + \dots \right) \right) \\ = 35$$

(which is achieved if all future developers develop their softwares and go open source). We can obtain a number of observations from these lower and upper bounds.

First, from the lower bound we can infer that $\sigma(10, 0, \cdot) \in \mathcal{O}$ (because the gross expected payoff of going open source is at least 20). This in turn implies that the gross expected payoff of going open source is strictly higher than 20, because the probability that the next developer developing his software is strictly positive.

Second, from the upper bound we can infer that $\sigma(40, 35, \mathbf{G}) = \mathbf{Q}$ and $\sigma(40, 35, \mathbf{B}) = \sigma(40, 35, \mathbf{R}) = \sigma(40, 35, \mathbf{1}) = \mathbf{P}$ (because by going open source the gross expected payoff is at most $35 < 50 = \pi_t + w_t$). It means the probability of any developer t

¹⁸Recall that if developer $t - 1$ did not choose any open source license, then either he did not develop software $t - 1$, or he did but chose to go proprietary. In either case, developer t would not have a chance to move.

¹⁹Recall that we can treat developer 0 as if he is subject to license \mathbf{B} .

going open source is at most $1 - P(40, 35) = (1 + \epsilon)/2 \approx 1/2$. This in turn implies a tighter upper bound for the gross expected payoff of going open source, which is approximately

$$W + \alpha_1 \left(\frac{W+w}{2} + \alpha_2 \left(\frac{W+w}{2} + \beta \frac{W+w}{2} + \beta^2 \frac{W+w}{2} + \dots \right) \right) = 33\frac{1}{8} < 35.$$

From this tighter upper bound we can infer that $\sigma(20, 20, \mathbf{B}) = \sigma(20, 20, \mathbf{R}) = \sigma(20, 20, \mathbf{1}) = \mathbf{P}$ (because by going open source the gross expected payoff is strictly less than $35 = \pi_t + w_t$).

Finally, since the gross expected payoff of going open source is strictly higher than 20, we have $\sigma(20, 20, \mathbf{G}) = \mathbf{G}$.

These observations imply that the equilibrium Markov strategy must be as depicted in Tables 7 and 8.

$\sigma(c, \pi, \mathbf{G})$	$\pi = 0$	$\pi = 20$	$\pi = 35$
$c = 10$	G	–	–
$c = 20$	–	G	–
$c = 40$	–	–	Q

Table 7: The equilibrium Markov strategy for a developer who is subject to **G**.

$\sigma(c, \pi, \mathbf{B/R/1})$	$\pi = 0$	$\pi = 20$	$\pi = 35$
$c = 10$	$o \in \mathcal{O}$	–	–
$c = 20$	–	P	–
$c = 40$	–	–	P

Table 8: The equilibrium Markov strategy for a developer who is subject to **B**, **R**, or **1**.

Now consider the problem of a developer t who is subject to license **B** and has drawn $(c_t, \pi_t) = (10, 0)$. If he chooses **G**, all future developers' behavior will be described by Table 7, and hence his gross expected payoff will be approximately

$$V_G = W + \frac{\alpha_1}{2} \left(W + \frac{\alpha_2}{2} \left(W + \frac{\beta}{2} W + \left(\frac{\beta}{2} \right)^2 W + \dots \right) \right) = 23\frac{1}{2}.$$

If he chooses either **B**, **R**, or **1**, with probability close to 1 the next developer will go proprietary, and hence his gross expected payoff is approximately

$$V_{B/R/1} = W + \alpha_1 w = 23\frac{3}{4} > 23\frac{1}{2} = V_G.$$

Therefore, $\sigma(10, 0, \mathbf{B}) \in \{\mathbf{B}, \mathbf{R}, \mathbf{1}\}$. To compare these three contenders, note that they differ only in developer t 's continuation payoff conditional on the event that $(c_{t+1}, \pi_{t+1}) = (10, 0)$, and hence it suffices to ask what developer t may want to allow developer $t + 1$ to choose in that event (hereafter event E).

Conditional on event E , if developer $t + 1$ chooses **G**, the present value of developer t 's continuation payoff will be approximately

$$V_{E:G} = \alpha_1 \left(W + \frac{\alpha_2}{2} \left(W + \frac{\beta}{2} W + \left(\frac{\beta}{2} \right)^2 W + \dots \right) \right) = \alpha_1 \times 28.$$

Conditional on event E , if developer $t + 1$ chooses either **B**, **R**, or **1**, the present value of developer t 's continuation payoff will be approximately

$$V_{E:B/R/1} = \alpha_1 (W + \alpha_2 w) = \alpha_1 \times 27\frac{1}{2} < \alpha \times 28 = V_{E:G}.$$

Therefore, when choosing among **B**, **R**, and **1**, developer t would like to choose the one that only allows developer $t + 1$ to choose **G**. This goal can be achieved by choosing **1**. We summarize these with the following proposition.

Proposition 2 *In the example in this subsection, a pure-strategy Markov perfect equilibrium always exists and is unique. In a pure-strategy Markov perfect equilibrium, developer 0*

1. *develops software 0 and goes proprietary when $(c_0, \pi_0) = (20, 20)$ or $(40, 35)$; and*
2. *develops software 0 and goes open source with the 1-chance-only license **1** when $(c_0, \pi_0) = (10, 0)$.*

What is going on? Intuitively, if developer $t + 2$ is allowed to go proprietary, he is more likely to develop software $t + 2$, but the stream of consumer surplus will also more likely terminate in period $t + 2$. To decide whether to allow developer $t + 2$ to go proprietary, one is hence trading off consumer surplus in period $t + 2$ versus consumer surplus in periods $s \geq t + 3$, which depends on that decision maker's discount rate between periods $t + 2$ and $t + 3$. With hyperbolic discounting, this discount rate is higher for developer $t + 1$ than for developer t , meaning that developer $t + 1$ is more tempted, compared to developer t , to allow developer $t + 2$ to go proprietary. Developer t hence may want to choose an open source license that forbids developer $t + 1$ from allowing developer $t + 2$ to go proprietary. While both **G** and **1** come with this same forbiddance, **1** has the extra benefit over **G** in allowing developer $t + 1$ to go proprietary. This is a benefit for developer t because, just like developer $t + 1$, developer t discounts the future hyperbolically and hence is equally tempted to allow the next developer to go proprietary.

The above intuition also explains why we need an $(\alpha_1, \alpha_2, \beta)$ -formulation to approximate hyperbolic discounting, instead of the more traditional (α, β) one. The key misalignment between developers t 's and $t + 1$'s interests lies in their discount rates between periods $t + 2$ and $t + 3$; i.e., between α_2 and β . The assumption that $\alpha_2 \leq \beta$ implies that developer t is more patient than developer $t + 1$, and hence will like to limit the latter's ability to reap an earlier reward. In contrast, α_1 does not play any role in this story, and the assumption that $\alpha_1 \leq \alpha_2$ can be relaxed without affecting our results.

8 Concluding Remarks

The open source movement is nothing short of a revolution in how production is organized. It has rightfully attracted a lot of economic studies, but very few on the very licenses that made this revolution possible. This paper is the first attempt to provide a theoretical framework to study all open source licenses, including

but not limited to the two most popular ones, namely GPL and BSD. We have provided an explanation of why GPL and BSD stood out from all these open source licenses as the most popular among open source developers. We have also presented settings where there are other open source licenses better than GPL and BSD.

One of the settings is where developers form a tree instead of a linear structure (see Subsection 7.1). Looking forward, we believe the most fertile ground for future research is to study how a developer can use different open source licenses to moderate the competition among his offsprings.

Studying a tree structure of developers, however, is exponentially more complicated than studying a linear structure of developers. Our example in Subsection 7.1 has already given a hint on one such complication, namely the necessity to specify the substitutability among softwares within the same generation. While the example in Subsection 7.1 involves only a minimal tree structure, the task of specifying substitutability can quickly become exponentially daunting in a more naturally-looking tree structure. For example, imagine a tree structure where every developer has two immediate offsprings: 1A and 1B from developer 0, 2Aa and 2Ab from developer 1A, 2Ba and 2Bb from developer 1B. . . , etc. It is natural to assume that software 2Aa is a closer substitute for software 2Ab than for software 2Ba. It means there are already two different levels of substitutability among generation-2 softwares. In general, there are at least n different levels of substitutability among generation- n softwares. It will need some clever modelling to keep such complications manageable without reducing the model back to one with a linear structure. We leave this difficult task for future research.

References

- [1] Athey, Susan and Glenn Ellison (2014), “Dynamics of Open Source Movements.” *Journal of Economics and Management Strategy*, 23(2): 294-316.

- [2] Balter, Ben (2015), "Open Source License Usage on GitHub.com." <https://github.blog/2015-03-09-open-source-license-usage-on-github-com/> (last accessed on November 23, 2022)
- [3] Brandenburger, Adam and Eddie Dekel (1993), "Hierarchies of Beliefs and Common Knowledge." *Journal of Economic Theory*, 59: 189-198.
- [4] Casadesus-Masanell, Ramon and Pankaj Ghemawat (2006), "Dynamic Mixed Duopoly: A Model Motivated by Linux vs. Windows." *Management Science*, 52(7): 1072-1084.
- [5] Economides, Nicholas and Evangelos Katsamakas (2006), "Two-Sided Competition of Proprietary vs. Open Source Technology Platforms and the Implications for the Software Industry." *Management Science*, 52(7): 1057-1071.
- [6] Fershtman, Chaim and Neil Gandal (2007), "Open Source Software: Motivation and Restrictive Licensing." *International Economics and Economic Policy*, 4: 209-225.
- [7] Fershtman, Chaim and Neil Gandal (2011), "A Brief Survey of the Economics of Open Source Software." In Palgrave Macmillan (ed.), *The New Palgrave Dictionary of Economics*, online edition.
- [8] Fielding, Roy T. (1999), "Shared Leadership in the Apache Project." *Communications of the ACM*, 42(4): 42-43.
- [9] Gaudoul, Alex (2004), "Open Source Software Development Patterns and License Terms." Toulouse working paper.
- [10] Gaudoul, Alex (2005), "Public Provision of a Private Good: What is the Point of the BSD License?" Toulouse working paper.
- [11] Gaudoul, Alex (2007), "Do Open Source Developers Respond to Competition? The (La)T_EX Case Study." *Review of Network Economics*, 6(2): 239-263.

- [12] Han, Xintong and Lei Xu (2019), "Technology Adoption in Input-Output Networks." Bank of Canada Staff Working Paper 2019-51.
- [13] Hertel, Guido, Sven Niedner, and Stefanie Herrmann (2003), "Motivation of Software Developers in Open Source Projects: An Internet-Based Survey of Contributors to the Linux Kernel." *Research Policy*, 32: 1159-1177.
- [14] Johnson, Justin P. (2002), "Open Source Software: Private Provision of a Public Good." *Journal of Economics and Management Strategy*, 11(4): 637-662.
- [15] Johnson, Justin P. (2006), "Collaboration, Peer Review and Open Source Software." *Information Economics and Policy*, 18: 477-497.
- [16] Lerner, Josh, Parag A. Pathak, and Jean Tirole (2006), "The Dynamics of Open-Source Contributors." *American Economic Review*, 96(2): 114-118.
- [17] Lerner, Josh and Jean Tirole (2002), "Some Simple Economics of Open Source." *Journal of Industrial Economics*, 50(2): 197-234.
- [18] Lerner, Josh and Jean Tirole (2005a), "The Economics of Technology Sharing: Open Source and Beyond." *Journal of Economic Perspectives*, 19(2): 99-120.
- [19] Lerner, Josh and Jean Tirole (2005b), "The Scope of Open Source Licensing." *Journal of Law, Economics, and Organization*, 21(1): 20-56.
- [20] Mariotti, Thomas, Martin Meier, and Michele Piccione (2005), "Hierarchies of Beliefs for Compact Possibility Models." *Journal of Mathematical Economics*, 41: 303-324.
- [21] Mertens, Jean-François and Shmuel Zamir (1985), "Formulation of Bayesian Analysis for Games with Incomplete Information." *International Journal of Game Theory*, 14: 1-29.

- [22] Mockus, Audris, Roy T. Fielding, and James D. Herbsleb (2002), "Two Case Studies of Open Source Software Development: Apache and Mozilla." *ACM Transactions on Software Engineering and Methodology*, 11(3): 309-346.
- [23] Raymond, Eric S. (1999), *The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary*, O'Reilly.
- [24] Roberts, Jeffrey A., Il-Horn Hann, and Sandra A. Slaughter (2006), "Understanding the Motivations, Participation, and Performance of Open Source Software Developers: A Longitudinal Study of the Apache Projects." *Management Science*, 52(7): 984-999.
- [25] Smith, P. McCoy (2022), "Copyright, Contract, and Licensing in Open Source." In Amanda Brock (ed.), *Open Source Law, Policy and Practice*, Oxford University Press.
- [26] Vendome, Christopher, Gabriele Bavota, Massimiliano Di Penta, Mario Linares-Vásquez, Daniel Germán, and Denys Poshyvanyk (2017), "License Usage and Changes: A Large-Scale Study on GitHub." *Empirical Software Engineering*, 22: 1537-1577.
- [27] von Hippel, Eric and Georg von Krogh (2003), "Open Source Software and the 'Private-Collective' Innovation Model: Issues for Organization Science." *Organization Science*, 14(2): 209-223.

Appendix A: GPL vs BSD

In this appendix, we shall demonstrate how, if we are willing to impose strong assumptions on the probability law $P(\cdot|\cdot)$, we can obtain some interesting characterizations on when developer 0 will choose GPL rather than BSD, and vice versa.

Consider the following example. Suppose $\beta = 1/2$; and $\forall t \geq 0, w_t \equiv w = 1$ and $W_t \equiv W = 3$. Suppose the joint distribution of c_t and π_t is iid across t , and is as depicted in Table 9, where z indexes the correlation between c_t and π_t (with $z = 1$ means perfect positive correlation, and $z = -1$ means perfect negative correlation).

$P(c_t, \pi_t)$	$\pi_t = 1$	$\pi_t = 7$
$c_t = 1$	$(1+z)/4$	$(1-z)/4$
$c_t = 7$	$(1-z)/4$	$(1+z)/4$

Table 9: Joint distribution of c_t and π_t .

Note that $c_t = 1$ or 7 with the same marginal probability, and $\pi_t = 1$ or 7 with the same marginal probability.

Since this is a stationary environment, we shall further restrict our attention to pure-strategy Markov perfect equilibrium. Our notations and definitions follow closely those in Subsection 7.2, with the exception that $\mathcal{O} = \{\mathbf{G}, \mathbf{B}\}$, $g(\mathbf{G}) = \{\mathbf{G}\}$, and $g(\mathbf{B}) = \{\mathbf{P}\} \cup \mathcal{O}$.

Consider any developer t . If he develops software t and goes open source, his gross payoff (gross of development cost c_t) is at least $W_t = W = 3$ (which is achieved if no future software is developed), and is at most

$$W_t + \beta W_{t+1} + \beta^2 W_{t+2} + \cdots = W/(1 - \beta) = 3/(1 - 1/2) = 6$$

(which is achieved if all future developers develop their softwares and go open source). Therefore,

1. if $c_t = 1$, he should develop software t , he should go open source if $\pi_t = 1$

- (because by going open source he can get at least $3 > 2 = \pi_t + w_t$), and he should go proprietary if $\pi_t = 7$ (provided he has such an option);
2. if $c_t = 7$, he should never go open source (because by going open source he can get at most 6), and he should not even develop software t if $\pi_t + w_t < 7$ (or, equivalently, if $\pi_t < 6$); and
 3. if $\pi_t = 7$ and he has the option to go proprietary, he should develop software t and then go proprietary.

These observations imply that the equilibrium strategy for a developer who is subject to license **G** (and hence cannot choose **P** or **B**) and for a developer who is subject to license **B**, respectively, must be the ones depicted in Tables 10 and 11.

$\sigma(c, \pi, \mathbf{G})$	$\pi = 1$	$\pi = 7$
$c = 1$	G	G
$c = 7$	Q	Q

Table 10: The equilibrium Markov strategy for a developer who is subject to **G**.

$\sigma(c, \pi, \mathbf{B})$	$\pi = 1$	$\pi = 7$
$c = 1$	$o \in \mathcal{O}$	P
$c = 7$	Q	P

Table 11: The equilibrium Markov strategy for a developer who is subject to **B**.

Now consider the problem of a developer t who is subject to license **B** when $(c_t, \pi_t) = (1, 1)$. If he chooses **G**, all future developers' behavior will be described by Table 10, and hence his gross expected payoff (gross of development cost) will be

$$V_G = W_t + (\beta/2)W_{t+1} + (\beta/2)^2W_{t+2} + \cdots = W/(1 - 1/4) = 3/(1 - 1/4) = 4.$$

If he chooses **B**, his payoff will depend on future developers' license decisions in the same situation; i.e., it will depend on $\sigma(1, 1, \mathbf{B})$. Suppose $\sigma(1, 1, \mathbf{B}) = \mathbf{B}$, then

his gross expected payoff will be

$$V_{BB} = W + \beta \left[\left(\frac{1+z}{4} \right) V_{BB} + \left(\frac{1-z}{4} + \frac{1+z}{4} \right) w \right] = \frac{26}{7-z}.$$

Suppose, instead, $\sigma(1, 1, \mathbf{B}) = \mathbf{G}$, then his gross expected payoff will be

$$V_{BG} = W + \beta \left[\left(\frac{1+z}{4} \right) V_G + \left(\frac{1-z}{4} + \frac{1+z}{4} \right) w \right] = \frac{15+2z}{4}.$$

An equilibrium with $\sigma(1, 1, \mathbf{B}) = \mathbf{B}$ exists iff $V_{BB} \geq V_G$; i.e., iff $z \geq 1/2$. An equilibrium with $\sigma(1, 1, \mathbf{B}) = \mathbf{G}$ exists iff $V_G \geq V_{BG}$; i.e., iff $z \leq 1/2$. We summarize these with the following proposition.

Proposition 3 *In the example in this appendix a pure-strategy Markov perfect equilibrium always exists, and is generically unique. In a pure-strategy Markov perfect equilibrium, developer 0*

1. *does not develop software 0 if development cost is high and potential profit is low (i.e., when $(c_0, \pi_0) = (7, 1)$);*
2. *develops software 0 and goes proprietary whenever potential profit is high (i.e., whenever $\pi_0 = 7$);*
3. *develops software 0 and goes open source when both development cost and potential profit are low (i.e., when $(c_0, \pi_0) = (1, 1)$); he goes open source using BSD if future development costs and potential profits are sufficiently positively correlated (i.e., if $z \geq 1/2$), and using GPL otherwise.*

Intuitively, developer 0 would like to see his idea further developed. However, development incurs development costs. In the unfortunate event that the next developer (i.e., developer 1) finds his development cost high, he will be discouraged from further developing developer 0's idea, especially if he is also prohibited from going proprietary and making a profit out of his effort. If development cost and

potential profit are positively correlated such that a higher potential profit typically accompanies a higher development cost, then BSD, by allowing developer 1 to go proprietary, will help encourage developer 1 to further develop developer 0's idea even in this unfortunate event. BSD hence can be the optimal choice for a developer going open source if development cost and potential profit are positively correlated, while GPL can be the optimal choice in the case of negative correlation.

Appendix B: The Universal Space of Open Source Licenses

In this appendix, we shall construct the universal space of open source licenses, $\mathcal{S}^U = (\mathcal{O}^U, g^U)$, that contains every space of open source licenses as its sub-space. The construction is similar to that of the universal type space in epistemic game theory.

We say that a space of open source licenses $\mathcal{S} = (\mathcal{O}, g)$ is *finite* if $g(o)$ is finite for every $o \in \mathcal{O}$. Note that the set of open source licenses \mathcal{O} needs not be finite even when the space $\mathcal{S} = (\mathcal{O}, g)$ is finite, as finiteness refers only to the number of options, $|g(o)|$, allowed by each open source license $o \in \mathcal{O}$.

To construct the universal space of open source licenses, $\mathcal{S}^U = (\mathcal{O}^U, g^U)$, we first recursively construct a sequence of nonempty sets $(\Theta_0, \Omega_0, \Theta_1, \Omega_1, \dots)$ as follows. Let $\Theta_0 = \Omega_0 = \{0, 1\}$. For $n \geq 1$, let

$$\begin{aligned}\Theta_n &= \mathcal{P}(\Omega_{n-1}) \\ \Omega_n &= \Omega_{n-1} \times \Theta_n \\ &= \Omega_{n-2} \times \Theta_{n-1} \times \Theta_n \\ &= \dots \\ &= \Omega_0 \times \Theta_1 \times \dots \times \Theta_n \\ &= \Theta_0 \times \Theta_1 \times \dots \times \Theta_n.\end{aligned}$$

A sequence $(\theta_0, \theta_1, \dots)$ is called a *restriction hierarchy* if $\theta_n \in \Theta_n$ for any $n \geq 0$. Let Ω_∞ be the set of all restriction hierarchies. Intuitively, an open source licenses can be represented by a restriction hierarchy $(\theta_0, \theta_1, \theta_2, \dots)$ with

- θ_0 specifying whether the next developer—say, developer t —is allowed to go proprietary (where $\theta_0 = 1$ means “yes” and $\theta_0 = 0$ means “no”),
- θ_1 specifying what restrictions developer t is allowed to impose on developer $t + 1$ regarding the option of going proprietary,

- θ_2 specifying what restrictions developer t is allowed to impose on developer $t + 1$ regarding *both* (i) the option of going proprietary *and* (ii) what restrictions developer $t + 1$ can impose on developer $t + 2$ regarding the option of going proprietary,
- \dots , etc.

To illustrate how an open source license can be represented by a restriction hierarchy, let's construct the restriction hierarchy $(\theta_0^G, \theta_1^G, \dots)$ that represents the GPL license **G**.

To determine θ_0^G , we ask whether GPL allows the next developer—say, developer t —to go proprietary. No, it does not. Therefore, $\theta_0^G = 0$.

To determine θ_1^G , we ask what options developer t has regarding whether to allow developer $t + 1$ to go proprietary. GPL gives developer t only a single option—meaning that θ_1^G must be a singleton. Moreover, that single option is to use GPL as well, whose first restriction has already been encoded in θ_0^G .²⁰ Therefore, θ_1^G must be the singleton $\{\theta_0^G\} = \{0\}$.

To determine θ_2^G , we ask what options developer t has regarding (i) whether to allow developer $t + 1$ to go proprietary, and (ii) whether to allow developer $t + 1$ to allow developer $t + 2$ to go proprietary. GPL gives developer t only a single option—meaning that θ_2^G must be a singleton. Moreover, that single option is to use GPL as well, whose first two restrictions have already been encoded in θ_0^G and θ_1^G .²¹ Therefore, θ_2^G must be the singleton $\{(\theta_0^G, \theta_1^G)\} = \{(0, \{0\})\}$.

More generally, for any $n \geq 1$, θ_n^G must be a singleton, and must be the singleton $\{(\theta_0^G, \theta_1^G, \dots, \theta_{n-1}^G)\}$.

The GPL license can hence be represented by the restriction hierarchy $(\theta_0^G, \theta_1^G, \dots)$,

²⁰The GPL license carries infinitely many restrictions. Here, we are concerned about only its first restriction, namely whether it allows the next developer to go proprietary.

²¹GPL carries infinitely many restrictions. Here, we are concerned about only its first two restrictions, namely (i) whether it allows the next developer to go proprietary, and (ii) whether it allows the next developer to allow the next-next developer to go proprietary.

where:

$$\begin{aligned}
\theta_0^G &= 0, \\
\theta_1^G &= \{0\}, \\
\theta_2^G &= \{(0, \{0\})\}, \\
\theta_3^G &= \{(0, \{0\}, \{(0, \{0\})\})\}, \\
\theta_4^G &= \{(0, \{0\}, \{(0, \{0\})\}, \{(0, \{0\}, \{(0, \{0\})\})\})\}, \\
&\vdots
\end{aligned} \tag{3}$$

Likewise we can construct the restriction hierarchy $(\theta_0^R, \theta_1^R, \dots)$ that represents the recursive-BSD license **R**. The recursive-BSD license allows the next developer to go proprietary, and hence $\theta_0^R = 1$. If the next developer chooses to go open source, the recursive-BSD license gives him only a single option—meaning that for any $n \geq 1$, θ_n^R must be a singleton. Moreover, that single option is to use the recursive-BSD license as well, whose first n restrictions have already been encoded in $\theta_0^R, \theta_1^R, \dots$, and θ_{n-1}^R . Therefore, θ_n^R must be the singleton $\{(\theta_0^R, \theta_1^R, \dots, \theta_{n-1}^R)\}$. The recursive-BSD license can hence be represented by the restriction hierarchy $(\theta_0^R, \theta_1^R, \dots)$, where:

$$\begin{aligned}
\theta_0^R &= 1, \\
\theta_1^R &= \{1\}, \\
\theta_2^R &= \{(1, \{1\})\}, \\
\theta_3^R &= \{(1, \{1\}, \{(1, \{1\})\})\}, \\
\theta_4^R &= \{(1, \{1\}, \{(1, \{1\})\}, \{(1, \{1\}, \{(1, \{1\})\})\})\}, \\
&\vdots
\end{aligned} \tag{4}$$

As our final example, let's also construct the restriction hierarchy $(\theta_0^1, \theta_1^1, \dots)$ that represents the 1-chance-only license **1**. The 1-chance-only license allows the

next developer to go proprietary, and hence $\theta_0^1 = 1$. If the next developer chooses to go open source, the 1-chance-only license gives him only a single option—meaning that for any $n \geq 1$, θ_n^1 must be a singleton. Moreover, that single option is to use GPL, whose first n restrictions have already been encoded in $\theta_0^G, \theta_1^G, \dots$, and θ_{n-1}^G . Therefore, θ_n^1 must be the singleton $\{(\theta_0^G, \theta_1^G, \dots, \theta_{n-1}^G)\}$. The 1-chance-only license can hence be represented by the restriction hierarchy $(\theta_0^1, \theta_1^1, \dots)$, where:

$$\begin{aligned}
\theta_0^1 &= 1, \\
\theta_1^1 &= \{0\}, \\
\theta_2^1 &= \{(0, \{0\})\}, \\
\theta_3^1 &= \{(0, \{0\}, \{(0, \{0\})\})\}, \\
\theta_4^1 &= \{(0, \{0\}, \{(0, \{0\})\}, \{(0, \{0\}, \{(0, \{0\})\})\})\}, \\
&\vdots
\end{aligned} \tag{5}$$

For any restriction hierarchy $(\theta_0, \theta_1, \dots)$, to the extent that it represents an open source license (not yet, as we shall see shortly), we should be able to extract from θ_0 whether it allows the next developer to go proprietary, and from $(\theta_1, \theta_2, \dots)$ what open source licenses it allows the next developer to use. The first task is easy ($\theta_0 = 1$ means “yes”, and $\theta_0 = 0$ means “no”), and the second task can be done by defining a correspondence $\Gamma : \Omega_\infty \rightrightarrows \Omega_\infty$ as follows: for any $(\theta_0, \theta_1, \dots) \in \Omega_\infty$ and $(\theta'_0, \theta'_1, \dots) \in \Omega_\infty$, $(\theta'_0, \theta'_1, \dots) \in \Gamma(\theta_0, \theta_1, \dots)$ iff $\forall n \geq 0$,²²

$$(\theta'_0, \dots, \theta'_n) \in \theta_{n+1} \in \Theta_{n+1} = \mathcal{P}(\Omega_n) = \mathcal{P}(\Theta_0 \times \Theta_1 \times \dots \times \Theta_n). \tag{6}$$

As an illustration, if we apply the correspondence Γ to restriction hierarchies

²²We abuse notation by writing $\Gamma((\theta_0, \theta_1, \dots))$ as $\Gamma(\theta_0, \theta_1, \dots)$, although Γ has only one argument, namely the vector $(\theta_0, \theta_1, \dots)$.

$(\theta_0^G, \theta_1^G, \dots)$, $(\theta_0^R, \theta_1^R, \dots)$, and $(\theta_0^1, \theta_1^1, \dots)$, we will obtain

$$\begin{aligned}\Gamma(\theta_0^G, \theta_1^G, \dots) &= \{(\theta_0^G, \theta_1^G, \dots)\}, \\ \Gamma(\theta_0^R, \theta_1^R, \dots) &= \{(\theta_0^R, \theta_1^R, \dots)\}, \quad \text{and} \\ \Gamma(\theta_0^1, \theta_1^1, \dots) &= \{(\theta_0^G, \theta_1^G, \dots)\},\end{aligned}$$

confirming the fact that a developer who is subject to open source license **G** (respectively, **R** and **1**), if going open source, is only allowed to do so using open source license **G** (respectively, **R** and **G**).

While an open source license can be represented by a restriction hierarchy, however, not every restriction hierarchy can be the representation of an open source license. In order for a restriction hierarchy $(\theta_0, \theta_1, \theta_2, \dots)$ to represent an open source license, it needs to satisfy a consistency requirement. To motivate this consistency requirement, note that both θ_1 and θ_2 contain information on what restrictions developer t is allowed to impose on developer $t + 1$ regarding the option of going proprietary, and these two pieces of information must agree with each other. More generally, for any $n \geq 1$, θ_n and θ_{n+1} must be consistent in the sense that $\text{Proj}_{\Omega_{n-1}} \theta_{n+1} = \theta_n$.²³

Definition 3 *A restriction hierarchy $(\theta_0, \theta_1, \dots)$ is consistent if $\text{Proj}_{\Omega_{n-1}} \theta_{n+1} = \theta_n$ for any $n \geq 1$.*

The reader can readily check that the restriction hierarchies $(\theta_0^G, \theta_1^G, \dots)$, $(\theta_0^R, \theta_1^R, \dots)$, and $(\theta_0^1, \theta_1^1, \dots)$ are all consistent, because for any $n \geq 1$,

$$\begin{aligned}\text{Proj}_{\Omega_{n-1}} \theta_{n+1}^G &= \text{Proj}_{\Theta_0 \times \dots \times \Theta_{n-1}} \{(\theta_0^G, \dots, \theta_n^G)\} = \{(\theta_0^G, \dots, \theta_{n-1}^G)\} = \theta_n^G, \\ \text{Proj}_{\Omega_{n-1}} \theta_{n+1}^R &= \text{Proj}_{\Theta_0 \times \dots \times \Theta_{n-1}} \{(\theta_0^R, \dots, \theta_n^R)\} = \{(\theta_0^R, \dots, \theta_{n-1}^R)\} = \theta_n^R, \quad \text{and} \\ \text{Proj}_{\Omega_{n-1}} \theta_{n+1}^1 &= \text{Proj}_{\Theta_0 \times \dots \times \Theta_{n-1}} \{(\theta_0^G, \dots, \theta_n^G)\} = \{(\theta_0^G, \dots, \theta_{n-1}^G)\} = \theta_n^1.\end{aligned}$$

²³ $\text{Proj}_{\Omega_{n-1}} \theta_{n+1}$ denotes the projection of $\theta_{n+1} \in \Theta_{n+1} = \mathcal{P}(\Omega_n) = \mathcal{P}(\Omega_{n-1} \times \Theta_n)$ into Ω_{n-1} , resulting in a subset of Ω_{n-1} . Consistency requires that this subset is the same as $\theta_n \in \mathcal{P}(\Omega_{n-1})$.

Let \mathcal{O}_1 be the set of all consistent restriction hierarchies.

Lemma 1 *For any consistent restriction hierarchy $(\theta_0, \theta_1, \dots) \in \mathcal{O}_1$, the subset $\Gamma(\theta_0, \theta_1, \dots) \subset \Omega_\infty$ is nonempty.*

PROOF: We first show that there exists $\theta'_0 \in \theta_1$. This follows from $\theta_1 \in \Theta_1 = \mathcal{P}(\Omega_0) = \mathcal{P}(\Theta_0)$ and $\mathcal{P}(\Theta_0)$ does not contain the empty set.

For any $n \geq 1$, assume that there exists $(\theta'_0, \dots, \theta'_{n-1}) \in \theta_n$. Since $(\theta_0, \theta_1, \dots)$ is consistent, we have $\text{Proj}_{\Omega_{n-1}} \theta_{n+1} = \theta_n$, and hence $(\theta'_0, \dots, \theta'_{n-1}) \in \text{Proj}_{\Omega_{n-1}} \theta_{n+1}$ as well. Therefore, there exists θ'_n such that $(\theta'_0, \dots, \theta'_{n-1}, \theta'_n) \in \theta_{n+1}$. By induction, there hence exists a restriction hierarchy $(\theta'_0, \theta'_1, \dots) \in \Omega_\infty$ such that, $\forall n \geq 0$, $(\theta'_0, \dots, \theta'_n) \in \theta_{n+1}$, implying that $\Gamma(\theta_0, \theta_1, \dots)$ is nonempty. ■

That the restriction hierarchy $(\theta_0, \theta_1, \dots)$ is consistent, however, does not guarantee that those restriction hierarchies in $\Gamma(\theta_0, \theta_1, \dots)$ are also consistent. Therefore, we shall define

$$\mathcal{O}_2 = \{(\theta_0, \theta_1, \dots) \in \mathcal{O}_1 : \Gamma(\theta_0, \theta_1, \dots) \subset \mathcal{O}_1\},$$

and retain only restriction hierarchies in \mathcal{O}_2 .

That every restriction hierarchy $(\theta'_0, \theta'_1, \dots)$ in $\Gamma(\theta_0, \theta_1, \dots)$ is consistent, however, does not guarantee that those restriction hierarchies in $\Gamma(\theta'_0, \theta'_1, \dots)$ are also consistent. Therefore, we shall not stop here, and shall continue and recursively define, for any $k \geq 2$,

$$\mathcal{O}_k = \{(\theta_0, \theta_1, \dots) \in \mathcal{O}_{k-1} : \Gamma(\theta_0, \theta_1, \dots) \subset \mathcal{O}_{k-1}\}.$$

We shall now define

$$\mathcal{O}^U = \bigcap_{k=1}^{\infty} \mathcal{O}_k,$$

which will be the set of every restriction hierarchy that represents an open source

license.²⁴

The reader can readily check that each of $(\theta_0^G, \theta_1^G, \dots)$, $(\theta_0^R, \theta_1^R, \dots)$, and $(\theta_0^1, \theta_1^1, \dots)$ belongs to O^U . Indeed, we have already seen that they are all consistent and hence belong to O_1 . For any $k \geq 1$, assume that they have already been shown to belong to O_k . Then we have

$$\Gamma(\theta_0^G, \theta_1^G, \dots) = \{(\theta_0^G, \theta_1^G, \dots)\} \subset O_k$$

by the inductive assumption, and hence $(\theta_0^G, \theta_1^G, \dots)$ belongs to O_{k+1} as well; and similarly for $(\theta_0^R, \theta_1^R, \dots)$ and $(\theta_0^1, \theta_1^1, \dots)$. Therefore, by induction, all of them belong to O_k for any $k \geq 1$, and hence belong to O^U . This also shows that O^U is nonempty.

Define $g^U : O^U \rightarrow \mathcal{P}(\{\mathbf{P}\} \cup O^U)$ such that, for any consistent restriction hierarchy $(\theta_0, \theta_1, \dots) \in O^U$,

$$g^U(\theta_0, \theta_1, \dots) = \begin{cases} \Gamma(\theta_0, \theta_1, \dots) & \text{if } \theta_0 = 0 \\ \Gamma(\theta_0, \theta_1, \dots) \cup \{\mathbf{P}\} & \text{if } \theta_0 = 1 \end{cases}. \quad (7)$$

We shall call the pair $\mathcal{S}^U = (O^U, g^U)$ the *universal space of open source licenses*. The following two theorems justify why this terminology is appropriate. The first theorem states that \mathcal{S}^U is in itself a space of open source licenses. The second theorem states that any finite irreducible space of open source licenses is a subspace of \mathcal{S}^U . These two theorems (except for the “1-to-1” part of Theorem 4) can also be proved as corollaries of Mariotti, Meier, and Piccione’s (2005) Proposition 3, by recognizing that any finite space of open source licenses can be made into a compact continuous possibility structure.²⁵ Below we provide an elementary proof without topology for each of these two theorems.

²⁴Readers familiar with the classical construction of the universal type space will recognize that this step is akin to the imposition of common knowledge of coherency.

²⁵We thank Yi-Chun Chen for pointing this out to us.

Theorem 3 *The universal space of open source licenses \mathcal{S}^U is in itself a space of open source licenses.*

PROOF: We have already seen that \mathcal{O}^U is nonempty, as it contains, for example, $(\theta_0^G, \theta_1^G, \dots)$, $(\theta_0^R, \theta_1^R, \dots)$, and $(\theta_0^1, \theta_1^1, \dots)$. For any $o \in \mathcal{O}^U \subseteq \mathcal{O}_1$, by Lemma 1, $\Gamma(o)$ is a nonempty subset of Ω_∞ . Moreover, for any $k \geq 1$, $\Gamma(o) \subseteq \mathcal{O}_k$ because $o \in \mathcal{O}^U \subseteq \mathcal{O}_{k+1}$, and hence $\Gamma(o)$ is also a nonempty subset of $\cap_k \mathcal{O}_k = \mathcal{O}^U$. Therefore, $g^U(o) \cap \mathcal{O}^U = \Gamma(o) \neq \emptyset$, and hence $\mathcal{S} = (\mathcal{O}^U, g^U)$ is in itself a space of open source licenses. ■

Theorem 4 *Any finite irreducible space of open source licenses $\mathcal{S} = (\mathcal{O}, g)$ is a sub-space of the universal space of open source licenses in the sense that there exists a 1-to-1 mapping $f : \mathcal{O} \rightarrow \mathcal{O}^U$, called the canonical representation, such that for any $o \in \mathcal{O}$,*

$$\begin{aligned} \mathbf{P} \in g(o) &\iff \mathbf{P} \in g^U(f(o)) \\ \text{and } o' \in g(o) &\iff f(o') \in g^U(f(o)). \end{aligned} \tag{8}$$

We prove Theorem 4 in three steps. We first explicitly construct a mapping f from \mathcal{O} into the set of all restriction hierarchies Ω_∞ . We call this mapping the *canonical representation* of open source licenses. We then prove that the range of f actually lies inside \mathcal{O}^U (Lemma 2). Finally, we prove that, if $\mathcal{S} = (\mathcal{O}, g)$ is finite irreducible, this will be a 1-to-1 mapping that satisfies (8) (Lemma 3).

We construct the canonical representation f by recursively defining a sequence of mappings (f_0, f_1, \dots) , with each f_n a mapping from \mathcal{O} into Θ_n . We first define f_0 such that, for any $o \in \mathcal{O}$,

$$f_0(o) = \begin{cases} 0 & \text{if } \mathbf{P} \notin g(o) \\ 1 & \text{if } \mathbf{P} \in g(o). \end{cases}$$

Suppose we have already defined mappings (f_0, \dots, f_n) , we then define f_{n+1}

such that, for any $o \in \mathcal{O}$,

$$f_{n+1}(o) = \{(f_0(o'), \dots, f_n(o')) : o' \in g(o) \cap \mathcal{O}\}.$$

We can now define the canonical representation $f : \mathcal{O} \rightarrow \Omega_\infty$ such that, for any $o \in \mathcal{O}$,

$$f(o) = (f_0(o), f_1(o), \dots).$$

Lemma 2 *The range of the canonical representation f lies inside \mathcal{O}^U .*

PROOF: We first prove that, for any $o \in \mathcal{O}$, $f(o)$ is a consistent restriction hierarchy. Let $f(o) = (\theta_0, \theta_1, \dots)$. For any $n \geq 1$,

$$\begin{aligned} \text{Proj}_{\Omega_{n-1}} \theta_{n+1} &= \text{Proj}_{\Omega_{n-1}} f_{n+1}(o) \\ &= \text{Proj}_{\Omega_{n-1}} \{(f_0(o'), \dots, f_n(o')) : o' \in g(o) \cap \mathcal{O}\} \\ &= \{(f_0(o'), \dots, f_{n-1}(o')) : o' \in g(o) \cap \mathcal{O}\} \\ &= f_n(o) = \theta_n, \end{aligned} \tag{9}$$

and hence $f(o)$ is a consistent restriction hierarchy.

Given any $o \in \mathcal{O}$. Consider a sequence $\{(\theta_0^k, \theta_1^k, \dots)\}_{k \geq 1}$ of restriction hierarchies such that

$$\begin{aligned} (\theta_0^1, \theta_1^1, \dots) &\in \Gamma(f(o)), \quad \text{and} \\ \forall k > 1, \quad (\theta_0^k, \theta_1^k, \dots) &\in \Gamma(\theta_0^{k-1}, \theta_1^{k-1}, \dots). \end{aligned} \tag{10}$$

We claim that, $\forall k \geq 1$, there exists $\{o_0^k, o_1^k, \dots\} \subset \mathcal{O}$ such that,

$$\forall n \geq 0, \quad (\theta_0^k, \dots, \theta_n^k) = (f_0(o_n^k), \dots, f_n(o_n^k)). \tag{11}$$

We prove this claim by induction. Consider $k = 1$. Since $(\theta_0^1, \theta_1^1, \dots) \in \Gamma(f(o))$,

we have, $\forall n \geq 0$,

$$(\theta_0^1, \dots, \theta_n^1) \in f_{n+1}(o) = \{(f_0(o'), \dots, f_n(o')) : o' \in g(o) \cap \mathcal{O}\},$$

and hence there exists $o_n^1 \in g(o) \cap \mathcal{O}$ such that $(\theta_0^1, \dots, \theta_n^1) = (f_0(o_n^1), \dots, f_n(o_n^1))$.

Assume we have already proved that, for some $k \geq 1$, there exists $\{\theta_0^k, \theta_1^k, \dots\} \subset \mathcal{O}$ that satisfies (11). Since $(\theta_0^{k+1}, \theta_1^{k+1}, \dots) \in \Gamma(\theta_0^k, \theta_1^k, \dots)$, we have, $\forall n \geq 0$,

$$(\theta_0^{k+1}, \dots, \theta_n^{k+1}) \in \theta_{n+1}^k = f_{n+1}(o_{n+1}^k) = \{(f_0(o'), \dots, f_n(o')) : o' \in g(o_{n+1}^k) \cap \mathcal{O}\},$$

and hence there exists $o_{n+1}^{k+1} \in g(o_{n+1}^k) \cap \mathcal{O}$ such that $(\theta_0^{k+1}, \dots, \theta_n^{k+1}) = (f_0(o_{n+1}^{k+1}), \dots, f_n(o_{n+1}^{k+1}))$.

By induction, the claim is hence true for all $k \geq 1$.

Given any $(\theta_0^k, \theta_1^k, \dots)$, for any $n \geq 0$,

$$\begin{aligned} \text{Proj}_{\Omega_{n-1}} \theta_{n+1}^k &= \text{Proj}_{\Omega_{n-1}} f_{n+1}(o_{n+1}^k) \\ &= f_n(o_{n+1}^k) = \theta_n^k, \end{aligned}$$

where the second equality follows from (9), and the last equality follows from $(\theta_0^k, \dots, \theta_n^k, \theta_{n+1}^k) = (f_0(o_{n+1}^k), \dots, f_n(o_{n+1}^k), f_{n+1}(o_{n+1}^k))$. Therefore, $(\theta_0^k, \theta_1^k, \dots)$ is a consistent restriction hierarchy.

Since this is true for any $(\theta_0^k, \theta_1^k, \dots)$ in any sequence $\{(\theta_0^k, \theta_1^k, \dots)\}_{k \geq 1}$ that satisfies (10), we have $f(o) \in \mathcal{O}_k$ for every $k \geq 1$, and hence $f(o) \in \mathcal{O}^U$. Since this is true for any $o \in \mathcal{O}$, we have proved that the range of f lies inside \mathcal{O}^U . ■

Lemma 3 *If $\mathcal{S} = (\mathcal{O}, g)$ is finite irreducible, the canonical representation f is a 1-to-1 mapping that satisfies (8).*

PROOF: Let $\mathcal{S} = (\mathcal{O}, g)$ be a finite irreducible space of open source licenses. Define an equivalence relation \sim on \mathcal{O} such that $\forall o_1, o_2 \in \mathcal{O}$, $o_1 \sim o_2$ iff $f(o_1) = f(o_2)$. Suppose, by way of contradiction, $f : \mathcal{O} \rightarrow \mathcal{O}^U$ is not 1-to-1, then \sim will be a

nontrivial equivalence relation. To arrive at a contradiction, it suffices to prove that \sim and g are compatible; i.e., it suffices to prove that, whenever $o_1 \sim o_2$,

$$(\mu \circ g)(o_1) = \{[x] : x \in g(o_1)\} =: A_1 = A_2 := \{[x] : x \in g(o_2)\} = (\mu \circ g)(o_2).$$

Since o_1 and o_2 play symmetric roles, it suffices to prove that, $[x] \in A_1 \implies [x] \in A_2$.

Suppose $[\mathbf{P}] \in A_1$, then $f_0(o_1) = 1$. But then $f_0(o_2) = 1$ as well because $f(o_1) = f(o_2)$, and hence $[\mathbf{P}] \in A_2$ as claimed.

Suppose $[o'] \in A_1 \setminus A_2$, then there exists $o'_1 \in g(o_1)$ such that $f(o'_1) = f(o')$, but for any $o'_2 \in g(o_2)$, $f(o'_2) \neq f(o')$. By finiteness of $g(o_2)$, there exists $n \geq 0$ such that

$$(f_0(o'), \dots, f_n(o')) \notin \{(f_0(o'_2), f_1(o'_2), \dots, f_n(o'_2)) : o'_2 \in g(o_2) \cap \mathcal{O}\}.$$

Since $o'_1 \in g(o_1)$, we have $(f_0(o'_1), \dots, f_n(o'_1)) \in f_{n+1}(o_1)$ by the construction of the canonical representation f . But then we have

$$\begin{aligned} (f_0(o'), \dots, f_n(o')) &= (f_0(o'_1), \dots, f_n(o'_1)) \\ &\in f_{n+1}(o_1) \\ &= f_{n+1}(o_2) \\ &= \{(f_0(o'_2), f_1(o'_2), \dots, f_n(o'_2)) : o'_2 \in g(o_2) \cap \mathcal{O}\}, \end{aligned}$$

a contradiction.

It remains to prove that $f : \mathcal{O} \rightarrow \mathcal{O}^U$ satisfies (8). By the construction of f and g^U , we have

$$\mathbf{P} \in g(o) \iff f_0(o) = 1 \iff \mathbf{P} \in g^U(f(o)) = \Gamma(f(o)) \cup \{\mathbf{P}\}.$$

Suppose $o' \in g(o)$. Then, by the construction of f , $\forall n \geq 0$, we have $(f_0(o'), \dots, f_n(o')) \in f_{n+1}(o)$. By the construction of Γ , we have $(f_0(o'), f_1(o'), \dots) \in \Gamma(f_0(o), f_1(o), \dots)$.

Therefore, by the construction of g^U , we have $f(o') \in g^U(f(o))$.

Suppose $f(o') \in g^U(f(o))$. Then, by the construction of g^U , we have $f(o') \in \Gamma(f(o))$. By the construction of Γ , $\forall n \geq 0$, we have $(f_0(o'), \dots, f_n(o')) \in f_{n+1}(o)$. By finiteness of $g(o)$, there exists $o'' \in g(o)$ such that $f(o'') = f(o')$. By the fact that f is a 1-to-1 mapping, we have $o' = o'' \in g(o)$. This completes the proof that f satisfies (8). ■

Appendix C: Categorizing Imposture-Proof Open Source Licenses

In this appendix, we shall categorize different imposture-proof open source licenses. This categorization is independent of our result in Section 6, which deals with the question of why GPL and BSD stood out from other licenses as the two most popular choices of open source developers. The categorization, however, may be useful for future research, especially in light of our results in Section 7, which demonstrate how open source licenses other than GPL and BSD can be useful if we go beyond our current model setup.

Given any space of open source licenses $\mathcal{S} = (\mathcal{O}, g)$ that is imposture-proof, and given any open source license $o \in \mathcal{O}$, let's recursively define the following:

$$\begin{aligned} G_0(o) &= g(o) \\ G_1(o) &= \bigcup_{o' \in G_0(o) \cap \mathcal{O}} g(o') \\ &\vdots \\ G_n(o) &= \bigcup_{o' \in G_{n-1}(o) \cap \mathcal{O}} g(o'). \end{aligned}$$

Since the space $\mathcal{S} = (\mathcal{O}, g)$ is imposture-proof, we must have

$$G_0(o) \supseteq G_1(o) \supseteq G_2(o) \supseteq \cdots .$$

Therefore, there exists a first time (possibly infinity), denoted by $L(o)$, such that the option \mathbf{P} is no longer available; i.e., $\mathbf{P} \in G_n$ iff $n < L(o)$. We say that $L(o)$ is the (*upper*) *level* of open source license o . Among the open source licenses studied in

Section 3, we have

$$\begin{aligned} L(\mathbf{G}) &= 0, \\ L(\mathbf{R}) &= \infty, \quad \text{and} \\ L(\mathbf{1}) &= 1. \end{aligned}$$

Lemma 4 *Let $\mathcal{S} = (\mathcal{O}, g)$ be an irreducible space of open source licenses that is imposture-proof. Any open source license $o \in \mathcal{O}$ with $L(o) = 1$ is identical to the 1-chance-only license $\mathbf{1}$ in the sense that $\exists \mathbf{G} \in \mathcal{O}$ with $g(\mathbf{G}) = \{\mathbf{G}\}$ such that $g(o) = \{\mathbf{P}, \mathbf{G}\}$.*

PROOF: By definition, $L(o) = 1$ implies that $\mathbf{P} \in G_0(o) = g(o)$ but $\mathbf{P} \notin G_1(o) = \cup_{o' \in g(o) \cap \mathcal{O}} g(o')$. Therefore, $\forall o' \in g(o) \cap \mathcal{O}$, $\mathbf{P} \notin g(o')$, and hence by Theorem 1 is identical to the GPL license in the sense that $g(o') = \{o'\}$. Irreducibility then implies that there is only one such o' , denoted by \mathbf{G} , and that $g(o) = \{\mathbf{P}, \mathbf{G}\}$. ■

For any $L \geq 1$, let's recursively define the L -chances-only license \mathbf{L} as follows.

L: *going open source with the L -chances-only license*

Developer t goes open source. If developer $t + 1$ chooses to develop software $t + 1$, he can either go proprietary (i.e., choosing \mathbf{P}), or go open source with either $\mathbf{G}, \mathbf{1}, \mathbf{2}, \dots, \mathbf{L-2}$, or $\mathbf{L-1}$.

Theorem 5 *Let $\mathcal{S} = (\mathcal{O}, g)$ be an irreducible space of open source licenses that is imposture-proof. Any open source license $o \in \mathcal{O}$ with $L(o) = L \geq 1$ is identical to the L -chances-only license \mathbf{L} in the sense that $\exists \mathbf{G}, \mathbf{1}, \mathbf{2}, \dots, \mathbf{L-1} \in \mathcal{O}$, with $g(\mathbf{G}) = \{\mathbf{G}\}$, $g(\mathbf{1}) = \{\mathbf{P}, \mathbf{G}\}$, $g(\mathbf{2}) = \{\mathbf{P}, \mathbf{G}, \mathbf{1}\}$, \dots , and $g(\mathbf{L-1}) = \{\mathbf{P}, \mathbf{G}, \mathbf{1}, \mathbf{2}, \dots, \mathbf{L-2}\}$, such that $g(o) = \{\mathbf{P}, \mathbf{G}, \mathbf{1}, \mathbf{2}, \dots, \mathbf{L-1}\}$.*

PROOF: By Lemma 4, the statement is true for $L = 1$. Assume we have already proved that the statement is true for $L = 1, \dots, m$, for some $m \geq 1$. Consider an open source license $o \in \mathcal{O}$ with $L(o) = m + 1$. By definition, $L(o) = m + 1$ implies

that $\mathbf{P} \in G_0(o), \dots, G_m(o)$ but $\mathbf{P} \notin G_{m+1}(o)$. Therefore, $\forall o' \in g(o) \cap \mathcal{O}, L(o') \leq m$, and $\exists o' \in g(o)$ such that $L(o') = m$. Therefore, by the inductive assumption there exist (and, by irreducibility, unique) $\mathbf{G}, \mathbf{1}, \mathbf{2}, \dots, \mathbf{m} \in \mathcal{O}$, with $g(\mathbf{G}) = \{\mathbf{G}\}$, $g(\mathbf{1}) = \{\mathbf{P}, \mathbf{G}\}$, \dots , and $g(\mathbf{m}) = \{\mathbf{P}, \mathbf{G}, \mathbf{1}, \dots, \mathbf{m} - \mathbf{1}\}$, such that

$$G_0(o) \subseteq \{\mathbf{P}, \mathbf{G}, \mathbf{1}, \dots, \mathbf{m}\} = g(\mathbf{m}) \cup \{\mathbf{m}\} \subseteq G_1(o) \cup \{\mathbf{m}\} \subseteq G_0(o),$$

and hence all inclusions are equality. By induction, the statement is true for any $L \geq 1$. \blacksquare

By imposture-proofness, if $L(o) = \infty$, then $\exists o' \in g(o) \cap \mathcal{O}$ such that $L(o') = \infty$, and hence $\max\{L(o') : o' \in g(o) \cap \mathcal{O}\} = \infty$. Let's define

$$l(o) = \min\{L(o') : o' \in g(o) \cap \mathcal{O}\}.$$

We shall call $l(o)$ the *lower level* of open source license o . Among the open source licenses studied in Section 3, we have

$$\begin{aligned} l(\mathbf{G}) &= 0, \\ l(\mathbf{R}) &= \infty, \quad \text{and} \\ l(\mathbf{1}) &= 0. \end{aligned}$$

Lemma 5 *Let $\mathcal{S} = (\mathcal{O}, g)$ be an irreducible space of open source licenses that is imposture-proof. For any open source license $o \in \mathcal{O}$, $l(o)$ is either 0 or ∞ .*

PROOF: Note that, for any open source license $o \in \mathcal{O}$, $\forall o' \in g(o), L(o') \geq l(o)$, and $\exists o' \in g(o)$ such that $L(o') = l(o)$. Suppose $l(o) = L < \infty$. Then, by Theorem 5, $\exists \mathbf{G}, \mathbf{1}, \mathbf{2}, \dots, \mathbf{L} \in \mathcal{O}$, with $g(\mathbf{G}) = \{\mathbf{G}\}$, $g(\mathbf{1}) = \{\mathbf{P}, \mathbf{G}\}$, $g(\mathbf{2}) = \{\mathbf{P}, \mathbf{G}, \mathbf{1}\}$, \dots , and $g(\mathbf{L}) = \{\mathbf{P}, \mathbf{G}, \mathbf{1}, \mathbf{2}, \dots, \mathbf{L} - \mathbf{1}\}$, such that $\mathbf{L} \in g(o)$. By imposture-proofness, $\mathbf{G} \in g(\mathbf{L}) \subset g(o)$, and hence $l(o) \leq L(\mathbf{G}) = 0$. \blacksquare

Theorem 6 *Let $\mathcal{S} = (\mathcal{O}, g)$ be an irreducible space of open source licenses that is imposture-proof. Any open source license $o \in \mathcal{O}$ with $L(o) = l(o) = \infty$ is identical to the recursive-BSD license \mathbf{R} in the sense that $g(o) = \{\mathbf{P}, o\}$.*

PROOF: Since $L(o) = \infty$, we have $\mathbf{P} \in g(o)$. If $o' \notin g(o)$ for any $o' \neq o$, then by the nonemptiness of $g(o) \cap \mathcal{O}$ we must have $g(o) = \{\mathbf{P}, o\}$, and we are done. Therefore, let's suppose there exists $o' \neq o$ such that $o' \in g(o) \cap \mathcal{O}$. Let's define an equivalence relation \sim such that $[o] = (g(o) \cap \mathcal{O}) \cup \{o\}$, and $[o''] = \{o''\}$ for any $o'' \notin (g(o) \cap \mathcal{O}) \cup \{o\}$. This is a nontrivial equivalence relation because $o' \neq o$ and yet $o' \sim o$. We shall prove that g and \sim are compatible, and hence $\mathcal{S} = (\mathcal{O}, g)$ is reducible. To prove compatibility, it suffices to prove that $(\mu \circ g)(o') = (\mu \circ g)(o)$ for any $o' \in g(o) \cap \mathcal{O}$. Since $l(o) = \infty$, we have $L(o') = \infty$, and hence $[\mathbf{P}]$ is contained in $(\mu \circ g)(o)$ as well as in $(\mu \circ g)(o')$. By imposture-proofness, we have

$$g(o') \cap \mathcal{O} \subseteq g(o) \cap \mathcal{O} \subseteq (g(o) \cap \mathcal{O}) \cup \{o\} = [o],$$

and hence $(\mu \circ g)(o') = \{[\mathbf{P}], [o]\} = (\mu \circ g)(o)$ as claimed. ■

We have, up to this point, characterized licenses $\mathbf{G}, \mathbf{1}, \dots, \mathbf{L}$, and \mathbf{R} . By Lemma 5, all the remaining imposture-proof open source licenses have the properties of $L(o) = \infty$ and $l(o) = 0$. These include the two BSD licenses studied in Sections 2 and 3, respectively, and many more.²⁶ Common across these licenses are:

1. All allows the next developer to go proprietary; i.e., $\mathbf{P} \in g(o)$.
2. All allows the next developer to go open source with a license $o' \in \mathcal{O}$ that is restrictive in the sense that $L(o') < \infty$.

²⁶Note that the two BSD licenses studied in Sections 2 and 3, respectively, are not exactly the same—while both have the flavor of “everything goes”, the meaning of “everything” differs across the two spaces of open source licenses studied in those two respective sections—and hence should more appropriately be regarded as different variants of BSD.

3. All allows the next developer to go open source with a license $o' \in \mathcal{O}$ that is permissive in the sense that $L(o') = \infty$.

Further categorization of these licenses will be left for future research.