

ONTOLOGY AND THE REGULATION OF INTELLECTUAL PROPERTY

1. Introduction

Philosophical reflection on intellectual property (IP) is still very young. Though lawyers have written much on the topic, the vast majority of this writing is philosophically unsophisticated. This paper aims to at least partially remedy this philosophical deficit by examining what reflection on the ontology of intellectual property can add to our understanding of how to regulate IP. I argue that ontological reflection should bring us to an important basic fact, namely that ownership of intellectual property involves the ownership of types rather than tokens. This difference in the ontological status of the objects owned makes a normative difference as to how we should regulate ownership of intellectual property as compared to tangible property.

I begin by arguing that the type-token distinction is the best way to account for the ontology of intellectual property. I next argue that the realisation that we are dealing with ownership of types rather than tokens has important normative implications. In particular some of the standard arguments in favour of private ownership of tangible property simply do not apply in the case of ownership of types, while others apply only in an attenuated way.

The following section examines the limitations of ontology as a guide to the regulation of IP. While thinking through the normative implications of the type-token distinction is a necessary condition for a sound regulatory approach to IP, it is certainly not sufficient. This is because many of the specific questions about the regulation of intellectual property that matter most in practice are not answerable by ontological reflection. Such questions concern the fair distribution of the burdens and benefits involved in the construction of public goods, rather than questions of ontology. I end by briefly sketching what an adequate normative regulatory model for IP might look like.

2. Towards an Ontology for Intellectual Property

The kinds of intellectual constructs that are owned and controlled through intellectual property regulations such as trademarks, patents, and copyrights are, even at a glance, different in kind from the things that are owned and controlled by laws which regulate the ownership of physical property such as cars and bicycles.

Cars and bicycles are rival goods: one person's use of them can interfere with another person's. For instance, you cannot ride your bike if I have borrowed it. Cars and bicycles are also subject to degradation through use: the more a car is driven, the quicker its parts will wear out. The types of intellectual constructs which are the subjects of intellectual property law are nonrival: one person's singing a song in no way interferes with another person's singing it, and similarly for use of computer software packages, and methods for making pharmaceuticals. Moreover, such intellectual constructs are not (except in exceptional circumstances) subject to degradation through overuse: as many people as they like can sing "Happy Birthday" without the song wearing out.¹ In Jefferson's words, ideas have the feature that "no one possesses the less, because every other possesses the whole of it. He who receives an idea from me, receives instruction himself without lessening mine; as he who lights his taper at mine, receives light without darkening me" (Jefferson 1813).

Ontology, at a very rough estimation, is the study of what there is. We can distinguish between two different ontological projects: descriptive ontology, and revisionary ontology (Strawson 1959, 10). Descriptive ontology attempts to clarify which kinds of entities are presupposed by our practices. Its aim is interpretive, and it does not attempt to say whether the types of entities that our practices presuppose really exist. Revisionary ontology attempts to get at the truth about what entities really exist, and aims to persuade us to revise our current ontology in line with the correct one. Little systematic philosophical thought has thus far been devoted to the ontology of intellectual property, and so it seems important at this stage to work out an adequate descriptive ontology which covers our current practices—even if we may in the fullness of philosophical reflection decide that the approach we should take to the ontology of intellectual property is a radically revisionary one.

It seems relatively clear that a descriptive ontology of our current IP practices should foreground the type-token distinction. The type-token

distinction seems to be the best way of interpreting our practices, in large part, because the other possibilities seem so obviously wrong. The relationship between the novel as owned and copyrighted by the author and individual copies of the novel is not a part-whole relationship, because all of the novel appears to be instantiated in each copy of it. Nor does the relationship between a property and its instances seem to fit the bill. The process as patented seems to be a structure that we deploy in creating instances of the machine which performs the process; it does not seem to be a mere way in which machines can differ.² Nor does a set-member relationship seem to be right, because sets have their members essentially, so that if another member is added to a set, or a member is taken away from it, we no longer have the same set. However, this is a very implausible way of thinking about novels and their copies: we do not think that anything metaphysically significant has happened to the novel *Great Expectations* itself if a copy of it is destroyed in a fire, or another is printed.

The type-token relation however seems to fit the bill appropriately, as types do not have the feature that they have to change when one of their tokens is destroyed or created, and they have a greater metaphysical solidity than properties:

The crucial difference between types and sets is this: while the identity of a set is determined by its actual membership, the identity of a type is determined not by which tokens actually exist but by the condition that something must meet to be one of its tokens. What makes the type K that type is that it lays down a certain condition for something to be one of its tokens; and it would still lay down this condition, and so would remain that type, even if fewer, more, or different tokens satisfied it. (Dodd 2008, 1119)

3. *Types, Tokens, and Ownership*

Ownership of tangible property involves ownership of tokens, whereas ownership of intellectual property involves ownership of types. However, most of the arguments in favour of private ownership revolve around factors that presuppose that ownership is of tokens rather than of types. Some of these arguments, of which we shall consider four, simply do not apply at all to types, while others have diminished force.

The first is Locke's argument for the necessity of private appropriation. Locke points out that if everything were held in common, and one had to get the consent of every other person before one were allowed to take anything from the common stock, then this would have absurd con-

sequences: everyone would be forced to either starve to death or to steal food from the common supply. Locke argues that given the importance of our right to self-preservation, it is legitimate to take private ownership of things which were previously held in common, so long as one does not allow what one has appropriated to go to waste, and one's act of appropriation leaves enough and as good for the next person (Locke 1690, II. 28).

While this argument seems powerful in the context of physical goods that are necessary for survival, the argument simply does not apply in the context of ownership of types (Shiffrin, 2001). Unlike tokens, types are nonrival; and a type can be tokened any number of times. So while my ability to read my token copy of Kant's first *Critique* is interfered with by your attempting to read my token copy, your reading a different token copy in no way interferes with my ability to read my token copy. So if this kind of argument from necessity or self-preservation has any force at all in the context of types, it seems to speak against ownership of types: if what matters is people being able to gain the access they need to essential goods which are necessary for their flourishing, then this access is hampered rather than facilitated by private ownership of intellectual property.

The second argument is the argument from the tragedy of the commons. There is reason to think that where there are resources which can be used by everyone and these resources are susceptible to degradation through overuse, these resources will tend to degrade over time. In Hardin's classic example, if we have a field that is held in common, then people will tend to overgraze it (Hardin, 1968). For it is in the interest of each shepherd to ensure that they have as many sheep as possible, and that each of their sheep are well-grazed; however if all (or most) shepherds behave in this way, then the commons will get overgrazed, and its ability to support sheep will soon be destroyed.³

However, the tragedy of the commons cannot occur in the case of types. Types are not worn out or degraded as more and more tokens of them are created; and so unlike a piece of land which can be overgrazed, there is no such possibility in the case of creating tokens of types. So while economic arguments about the degradation of the commons have frequently been used as arguments in favour of regimes of private ownership of real property, the same arguments cannot support private ownership of ideas.⁴

The third and fourth arguments have some force in the case of ownership of types, but this force is diminished from their force in the case of ownership of tokens. The third argument is the argument from fair

benefit, where the basic idea is that when a person labours to create a new object or to improve an otherwise unowned or commonly owned object, then their labouring gives them a claim to the exclusive use of the resource that they have improved or created.

Many if not most people feel the intuitive force of this idea. However there are two different readings of the underlying moral principle that justifies this judgement and only one of these readings gives strong support to ownership of types. The first reading would be that if someone, through their own labour, improves a resource that is unowned or held in common, then they thereby gain a right to the full use of that item—and owing to the fact that such goods are rival and that others' use will tend to limit the labourer's ability to use the item they have improved or created—they thereby also have the right to exclude others from its use. On the second reading, if someone through their own labour improves a resource that is unowned or held in common, then they thereby gain a right to the full use of that item and the right to exclude others from its use, even where others' use might not limit their use of it.

On the first reading, the right to exclude others only accrues to the person who labours because of the type of thing to which the person has laboured: because tokens are rival, a person cannot gain full ownership and use of a token without being able to exclude others from the use of that token. If we adopt this reading then the right to exclude would not carry over to the case of types: for the person who devises a new dance step could dance their new step as much as they like, regardless of whether others are also dancing the step. It is only if we adopt the second reading that it would follow that labouring to create or improve a type should give you a right to exclude others from the nonrivalrous use of that type. Which reading is to be preferred? I am inclined to think that the first reading better expresses our considered moral convictions, because it is far from clear to me why we should think that anyone has a morally compelling interest in excluding others from nonrivalrous uses of types.⁵

The fourth argument is the argument from productivity. The basic idea of this argument is that unless we allow private appropriation, there will be insufficient incentive for people to put in the necessary effort to improve resources that are currently unowned or held in common. If everyone is able to benefit from my work, then it will typically be irrational in self-interested terms for me to put in the effort to improve the resource unless I can exclude others from its use.

This argument does have some force for the ownership of types, but again we need to be careful in thinking about the differences we introduce in switching from tokens to types. Types are by nature much more difficult to exclude other people from than are tokens (a particular piece of land can be fenced off, and a bicycle can be locked away in a shed, but it is much more difficult to make ideas exclusive). So to that extent, it might look as if there is an even greater need for private intellectual property to ensure productivity than there is for private physical property. However, the nonrivalrousness of types means that nonexclusivity is in a different sense less of a problem than it is in the case of rival goods. If I have sufficient self-interested reason to work to improve or devise a type for my own purposes, then the fact that the type is nonexclusive will not dissuade me from putting in the effort. For instance, many of the people who contribute to open source software describe their motivation in self-interested terms: they have a particular problem they want to solve, and they write a program to solve it. The fact that others can then benefit from their work need not provide a disincentive for them, and if they are even moderately altruistic may provide an incentive. This makes the case of software rather different from, say, growing tomatoes: my ability to enjoy my tomatoes is interfered with if you eat them; while my ability to use my program is not interfered with if you use it too.

Taken together these ontological reflections show that private intellectual property is rather more difficult to justify than private physical property. Certainly this ontological reflection should bring us to see that there cannot be any pre-legislative entitlements to ownership of private intellectual property. As Jefferson puts it, "Inventions then cannot, in nature, be a subject of property. Society may give an exclusive right to the profits arising from them, as an encouragement to men to pursue ideas which may produce utility, but this may or may not be done, according to the will and convenience of the society, without claim or complaint from anybody" (Jefferson 1813; also Wilson 2009).

4. The Limitations of Ontology

Recognising that ownership of types has different normative implications from ownership of tokens is vital in thinking through how we should regulate intellectual property. But there are at least four ways in which our current systems for intellectual property regulation depart from the bare type-token distinction. Given this, the bare type-token distinction

is not sufficiently conceptually rich to explain how we currently regulate IP, and it would require a rather revisionary (and as we shall see, unwisely revisionary) approach to bring our current practices in line with the bare type-token distinction. Ontology can do no more than provide a rough guide, and cannot on its own tell us how to regulate IP. We need a richer approach to IP, which takes to heart the lessons we have just seen about the ownership of types, but places them in a broader regulatory context in which egalitarian considerations of fairness and requirements of incentivisation are reconciled.

The first and most obvious difference between the type-token relationship and our current regime for the regulation of IP is that intellectual property regulations are of limited duration, while type-token relationships are not. If something has the right features to be a token of type X, then anything in the future with the same features will also be a token of the same type. However judgements of infringement in intellectual property depend not just on the question of whether the new object has the right features to be a token of the type X, but also whether the type X is still within the duration of the relevant intellectual property regulations. This means that performing one and the same action, such as making use of a certain formula to manufacture a pharmaceutical, will have a different legal status depending on whether the drug is still within the patent period. It will be an infringement of intellectual property rights to make a generic version of a pharmaceutical one year and not the next, though the law has not changed.

Someone might take the view that the ontological distinction between type and token is the only thing that matters in IP regulation. On such a view, we should strike down all the elements in IP regulation that are not derivable from the ontology of the type-token distinction, and hence we should not make our intellectual property regulations time-bound. However, this would be a deep mistake. Regulation is properly responsive to human interests and to reflection on the type of society we are trying to create; while it should not incorporate any obvious ontological errors, it is a mistake to think that we should not include any distinctions in our regulations for which we cannot find corresponding ontological differences.⁶ Moreover, the ontology of intellectual property shows that it is more difficult to justify private ownership of intellectual property than private ownership of physical property. Hence it is better to think of the ontology here as setting limits to the kind of ownership claims that are legitimate,

rather than licensing the claim that because the type-token relationship is eternal, so should be our intellectual property restrictions.

The second way in which our current intellectual property regime is more complex than the bare type-token distinction relates to causation. The type-token distinction considered in itself leaves it open whether a token has to have a particular causal history in order to count as a token of a particular type. (Suppose, for instance that a piece of metal physically indistinguishable from a one Euro coin arises spontaneously through a natural process. Does this count as a token of the type of one Euro coin?) Under our current regime, some kinds of intellectual property rights protect against the creation of tokens only with a certain causal history, while others protect against the creation of new tokens regardless of their causal history. Copyright and trade secrets prohibit copying of the protected ideas, but do not make it illegal to invent the same object independently. In the case of patents, independent invention is no defence (see Trerise, this volume).

The third difference relates to novelty. Looking from the direction of ontology, there is no requirement that each type be in any way original, or more than minimally different from other types. Copyright follows this approach, with a very low requirement for originality before a copyright can be claimed, but we require that a patent shows a novel and non-obvious inventive step, so that, in order to count as patentable, a type has to be sufficiently far removed from already known types.

Putting these two differences together, we see that copyright imposes a low bar on what you have to do in order to be able to claim a copyright, whereas patent with its requirement of an inventive step implies a much higher bar. Copyright, however only covers creation of ideas that are influenced by the thing copyrighted, whereas patent covers also inventions which incorporate the idea but which were invented entirely separately from it.

Again, it seems implausible to think that we could find sufficient ontological differences to explain these different ways of regulating the production of ideas in different domains. A revisionary view might say that we should align all our intellectual property regulation with the type-token distinction. But again there seems little reason to do so. The various ways that the type-token distinction is interpreted in the different branches of IP regulation seem to reflect judgements about the nature of the incen-

tives that it is necessary to provide in different areas, and judgements about what is valuable about different types of creative activity. These are precisely the kinds of ethical judgements we need if we are to regulate in a fair and effective manner.

The fourth point of difference concerns the specificity of types. The type-token distinction leaves open the possibility (and is usually thought to imply) that there can be an infinite number of types, and that one and the same object can stand in the relationship of token to any number of very slightly different types. For example, one and the same chair can be a token of the type of furniture, the type of items designed by Bernard Schmidt, the type of items in the IKEA catalogue, and so on. This fecundity of types and tokens threatens to be problematic in the case of intellectual property regulation. For in judging whether one work infringes (or ought to be taken to infringe) the copyright on another, we want to know not whether there is a type such that the two pieces of writing are both tokens of that type (presumably there will always be a sufficiently general type such that any two things count as tokens of it), but rather whether both pieces should count as tokens of the *relevant* type.

We could if we wished explicate the relevant type in wider or in narrower ways. For instance, at its very narrowest, we could identify the type with a particular sequence of words. In this case, changing a single word (let alone translating the work into a different language) would mean that we were no longer dealing with tokens of the same type. Or at its broadest we could think of the relevant type in terms, say, of underlying plot structures (in which case we could say that all romantic comedies would fall under the same type).⁷

Could further ontological reflection help us to pick out what we should take the relevant type to be? I do not think so, because it is not the case that the readings that make the relevant type either narrower or broader than those we ourselves favour fail to pick out a genuine type at all. For instance, it seems that the particular sequence of words that Eliot published as *The Wasteland* is a genuine type, of which there can be tokens. The problem with taking this sequence to be the relevant type would not be ontological: rather it would be that there are reasons internal to the goals of what we are trying to achieve in regulating intellectual property which make it unhelpful to pick out the bare sequence of words as the relevant type. Picking such a narrowly defined type would create

problematic loopholes in intellectual property regulation, and partly because of this would tend to provide too small an incentive to creators to produce new work.⁸

In short, the central question we need to answer in regulation of intellectual property (and which will require differential answers in different domains) is: “when should a particular use of ideas count as an infringement of the relevant intellectual property regulation?” Ontological investigation can suggest a formal answer, namely “when it is an unauthorised tokening of a relevantly protected type”, but ontology cannot determine what the relevant types are.

5. Conclusion: Regulation Beyond Ontology

In closing I shall indicate very briefly how we might integrate the normative lessons from the ontology of intellectual property into a coherent regulatory strategy (for details, see Wilson forthcoming). I suggest that we should begin by making a distinction between rights and social goals. Rights are legitimate moral claims that individuals have and that ought (in usual circumstances) to be granted priority over other legitimate goals.⁹ Rights thus defined are highly resistant to aggregation: we cannot (in general) violate one person’s rights in order to prevent others having their rights violated. Each rights holder has a separate and weighty complaint if their rights are violated. Goals are states of affairs that governments should aim to promote.¹⁰ Goals are much less resistant to aggregation than are rights: governments will often have to make tradeoffs in the pursuit of their goals, and—assuming the tradeoffs are performed in a nondiscriminatory and reasonable way—this is perfectly legitimate. Unlike rights, individuals do not have separate and morally weighty complaints where their government is making good-faith attempts to pursue a goal, but—due to reasonable tradeoffs with the pursuit of other normatively valuable goals—the government has not yet provided the aimed-for level of service to that person.¹¹

Given this distinction, the first task for regulators is to ensure that they do not violate rights in the ways they regulate or fail to regulate. Hence, when we are thinking about the regulation of IP we should start by thinking about rights, and look and see whether there are any ways of regulating (or refusing to regulate) IP that might violate rights. There are two obvious ways in which the regulation of IP might violate rights. First,

refusing to enact robust private ownership rights over intellectual property might violate the rights of creators; and second, enacting robust private intellectual property legislation might violate the rights of those denied the ability to make use of these privately owned objects.

I have argued that there are various features of types as opposed to tokens that make it more difficult to justify private ownership in the case of types than of tokens, and I have further argued that this should lead us to think that there are no rights in this strong sense of rights to own intellectual property. A further question arises as to whether it is possible to violate the rights of potential users of intellectual property by depriving them of access to IP goods. If there are positive rights to basic medicines, then it seems plausible to think that in certain circumstances we may violate these rights.¹²

The second layer is that of social goals. The normative goal served by private intellectual property is the production of an adequate supply of intellectual goods. Private ownership of IP is justifiable—where it is—because it is, all things considered, the best solution to the public goods problem in the regulation of the production of new ideas. There are other normative goals, such as the construction and maintenance of a society of equals, which will tend to conflict with the use of private intellectual property as a strategy for ensuring an adequate supply of intellectual property goods. A society of equals is one in which each citizen can look each other in the eye and think of herself as of equal status to each other person. The goal of a society of equals is undermined where there are goods which have a large effect on the way which social status is negotiated, and which are differentially spread (particularly when this reinforces existing patterns of advantage and disadvantage). Conversely such a society is promoted where goods and freedoms that are important for social status are available to each on conditions of equality.

Nonrival goods are not capable of scarcity, and are hence capable of being supplied to everyone who desires them. If there is a scarcity in the supply of a given nonrival good, it is because we have elected to create an artificial scarcity. This is different from the case of rival goods where we frequently see “natural” scarcities. Because of this, ideas as goods have a particular force for the goal of achieving a society of equals: because ideas matter for human life, and because ideas can be made available to all at only a marginal extra cost, the goal of a society of equals will push us towards open access to ideas.

Working out how to balance our commitments to equality and the need to provide sufficient incentives to creators to ensure an adequate supply of intellectual property goods is a complex task that I cannot undertake here. The main point of this article has been to argue that reflection on the ontology of intellectual property should form an essential part of the background to our regulatory framework for IP, but that working out the specifics will require further first-order normative thinking.

James Wilson

*Centre for Philosophy, Justice, and Health
University College, London*

NOTES

1. Some information, such as racing tips, is rival: the more people who know that Baskettease has a good chance of winning the 3.15, the less favourable the odds will become. In addition Landes and Posner (2003) argue that works released into the public domain may suffer from congestion externalities, and that if so, the value of the underlying work may be better husbanded by restricting access to it rather than allowing free access. Their examples involve cases where a trademark (such as Mickey Mouse) could reduce in value because of brand dilution.

2. The ontology of intellectual property objects is similar in many ways to the ontology of musical works. Dodd argues on the same grounds that we should not think of the ontology of music in terms of properties and instances: “While a property is repeatable inasmuch as it is an entity capable of multiple instantiation by particulars, Vaughan Williams’s *Oboe Concerto* looks categorially unsuited to be identified with a property of its performances and playings. Rather than being a mere respect in which performances or playings can be alike or differ, the work itself is the blueprint for such performances and playings: a thing in its own right.” (Dodd 2008, 1118)

3. It is worth pointing out—as Ostrom (1991) and others have argued—that even where rival goods are held in common, the tendency towards a tragedy of the commons is by no means inevitable; rather there are various ways of regulating the commons which can successfully protect and sustain it. So even in the case of common ownership of physical resources, it is unclear how seriously we should take the idea of the tragedy of the commons as an argument in favour of private property.

4. Insofar as there are sound economic arguments for intellectual property they will have to focus on the incentives for the production of such goods, not the fear that they will otherwise be used up, as in argument four below.

5. As Penner puts it, “The right to property is grounded by the interest we have in using things in the broader sense. No one has any interest in merely excluding others from things, for any reason or no reason at all” (2000, 70). I argue for these claims at greater length in Wilson (2009), and also attempt a diagnosis of how and why some people nonetheless find the second reading more plausible.

6. To give just one example, many countries legislate to ensure that buildings of special architectural interest are protected against arbitrary changes by their owners. Such a way of regulating is legitimate despite the fact that it seems very implausible to think that there are sufficient ontological differences between architecturally uninteresting and architecturally interesting buildings to justify this.

7. Over the past 150 years there has been a shift in copyright law from narrower towards broader readings of what the relevant type is. For instance, when Harriet Beecher Stowe tried to prevent the publication of unauthorised translations of *Uncle Tom's Cabin*, she lost on the grounds that copyright covered only copying of the relevant text and not its translation. For a good overview of some of these changes, see Fisher (1999).

8. Adopting an over-broad conception of the relevant type will give us opposite problems: it is likely to hamper the efforts of future creators so much that we will judge it would have been better to have provided a narrower band of protection.

9. See for example Dworkin: "Rights are best understood as trumps over some background justification for political decisions that states a goal for the community as a whole. If someone has a right to publish pornography, this means that it is for some reason wrong for officials to act in violation of that right, even if they (correctly) believe that the community as a whole would be better off if they did" (1984, 165–66).

10. We can distinguish between de facto goals (what states of affairs a government is in fact aiming to promote), and normative goals (what states of affairs a government ought to be aiming to promote). Our interest is in normative goals and the regulation of intellectual property.

11. For instance, if you treat the desideratum that a person not wait for more than three months before having a particular type of operation as their right, then that person's right has been violated if they have to wait four months before having their operation. If however we treat this desideratum merely as a goal, then there is no reason to think that the person who has to wait for four months has been wronged, so long as the government has been making good-faith approaches towards the goal.

12. And if Pogge (2006) is correct, intellectual property rights may be able to violate even negative rights.

REFERENCES

- Dodd, Julian 2008. "Musical Works: Ontology and Meta-Ontology," *Philosophy Compass* 3(6): 1113–34.
- Dworkin, Ronald 1984. "Rights as Trumps" in Jeremy Waldron, ed., *Theories of Rights*, Oxford University Press: Oxford, 153–67.
- Fisher, William 1999. "The Growth of Intellectual Property: A History of the Ownership of Ideas in the United States," *Eigentumskulturen im Vergleich*, Vandenhoeck & Ruprecht, 265–91.
- Hardin, Garrett 1968. "The Tragedy of the Commons," *Science*, 162: 1243–48.
- Jefferson, Thomas 1813. Letter to Isaac McPherson, 13 August 1813, in *Writings*, 13: 333–35.
- Landes, William and Richard Posner 2003. *The Economic Structure of Intellectual Property Law*, Cambridge, Mass: Harvard University Press.

- Locke, John 1690. *Two Treatises of Government*, ed. Peter Laslett, Cambridge: Cambridge University Press.
- Ostrom, Elinor 1991. *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge: Cambridge University Press.
- Pogge, Thomas 2006. "Montréal Statement on the Human Right to Essential Medicines," *Cambridge Quarterly of Healthcare Ethics*, 15(2).
- Shiffrin, Seana 2001. "Lockean Arguments for Private Intellectual Property," in Stephen Munzer, ed., *New Essays in the Legal and Political Theory of Property*, Cambridge: Cambridge University Press, 138–67.
- Strawson, Peter 1959. *Individuals*, London: Routledge.
- Wilson, James 2009. "Could There Be A Right to Own Intellectual Property?" *Law and Philosophy*, 28(4): 393–427
- . (forthcoming). "On the Value of the Intellectual Commons," in Annabelle Lever, ed., *Philosophy and Intellectual Property*.