


**On the Very Idea of Owning Ideas:
Philosophical Reflections on Intellectual Property**

Bryan Cwik
Chicago, Illinois

BA, Fordham University, 2004
MA, University of Arkansas, 2007

A Dissertation presented to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of
Doctor of Philosophy

Corcoran Department of Philosophy
August 2013



My dissertation examines a number of philosophical questions about intellectual property (or IP). Within the broad conversation about IP going on in multiple disciplines, a number of philosophical claims and questions are floating around, and all too often are not subjected to the proper level of analysis. I examine, in a series of discrete but related chapters, some of the central positions about IP, with the aim of deriving general ideas about it from these ground-level discussions.

Three themes run throughout the essay. The first is a general opposition to instrumentalism about IP. Instrumentalism holds that most of the significant questions about IP can be reduced to questions about whether IP is required to incentivize intellectual labor, and thus ensure an adequate supply of creative, artistic, and innovative work. The second theme is the gap between a defensible conception of IP and the institutions that we currently have. Though the essay has a definite pro-IP slant, the kinds of institutions discussed and defended here are very different from IP institutions as they currently are, and it thus presents a basis for *criticism* of existing IP institutions. The third theme is a criticism of framing IP policy questions as tradeoffs between the costs of IP and the goal of ensuring an adequate supply of intellectual labor.

I argue here that settling the central questions about IP involves deep normative issues, about the kind of creative and innovative work we want, and the role of government and law in regulating these endeavors. This resists a crude “cost/benefit” framework, and requires us instead to pay attention to the kinds of goods incentivized by IP institutions, the different individuals affected by IP, and the constraints (such as demands of distributive justice) these institutions must operate within. The essay aims to demonstrate how these considerations can be incorporated into IP theory and policy. Its goal is thus a broadening and deepening of the multidisciplinary conversation about IP, and the role of these institutions in shaping and regulating art, science, culture, and technology.

For all my teachers, past, present, and future,
especially Eileen McCardle, Maggie Stewart,
Mel Graham, Jim Marsh, Mike Latham,
Jack Lyons, Sahar Akhtar, John Arras,
and – *most* especially – Loren Lomasky

Preface and Acknowledgments

In a report to the Congressional Subcommittee on Patents, Trademarks, and Copyrights in 1958, the eminent economist Fritz Machlup remarked that, if we did not have a patent system, he could not recommend instituting one; however, as we already have one, he could not in good conscience recommend abolishing it (Machlup 1958). Machlup's remark expresses the perfect attitude towards the tangle of theoretical and policy questions we confront about intellectual property. Though it may seem as if Machlup was simply ambivalent, or perhaps just lacked the intellectual courage to tell Congress his real opinion, I think the remark shows that Machlup appreciated that intellectual property was a topic of immense complexity, and significant issues about it resisted reduction to questions about "having or not having" patents and copyrights.

This essay is offered in a spirit of Machlupian care. I consider, in discrete yet related chapters, some of the major arguments and issues involving intellectual property from both technical and popular discussions. These chapters can be read independently of each other, but they are related by a few general themes. The central theme is a rejection of reducing issues about intellectual property to a single question, about the necessity of patents and copyrights for providing incentives to artists and inventors, or about the tradeoff between costs and benefits of intellectual property, or about whether there are sufficient "moral" reasons in favor of these institutions to justify them, regardless of concerns about incentives or net social benefit. I hope to show that this sort of reductionism simply will not work, and to illuminate the genuinely *philosophical* dimension of all the difficulties surrounding intellectual property.

The idea of writing something about intellectual property came to me in the spring of 2009, when I began reading some of the work of James Boyle. A conversation with my sister, Marissa Cwik (at the time a law student at Vanderbilt) during one of her visits to Charlottesville convinced me that it was worth thinking about, but I never considered intellectual property a potential *dissertation* topic until the following spring, when Loren Lomasky suggested it to me as such. After a few weeks of tinkering, I decided that it would do just fine as a dissertation project, and I announced my intention at a fateful dinner with Loren, Geoffrey Brennan, and Geoffrey Sayre-McCord. Their enthusiastic endorsement of the idea, sealed with a toast, was the all-important first sign that writing on such an unusual topic (for a philosophy dissertation) would not be a big mistake.

During the writing of this dissertation, I received financial support from the University of Virginia philosophy department in the form of an “in-house” dissertation year fellowship, and from an Adam Smith Fellowship from George Mason University’s Mercatus Center. Chapter 3 was presented at an Institute for Humane Studies Summer Seminar at Towson University in June 2012. A version of that chapter, under the same title, is forthcoming in *Ethical Theory and Moral Practice*. A distant ancestor of Chapter 4 was presented at a “Philosophy After Dark” colloquium to the UVa philosophy department in September 2012. Chapter 5 was presented at Penn State University in March 2013, and at a Mercatus Center Colloquium at George Mason University in June 2013. Thanks to the audiences at those talks, and two anonymous referees for *Ethical Theory and Moral Practice*, for their comments.

Over the years I wrote this essay I incurred many debts, both professional and

personal. My greatest (of both kinds) is of course to my dissertation supervisor, Loren Lomasky. Loren not only encouraged my work on this topic and read my drafts assiduously, he also pushed me hard and stuck to the deadlines we set, and let me develop my views on the subject largely without interference – both rare qualities in an advisor. Loren is not only my supervisor, he is also the best teacher I’ve ever had, and I count myself very fortunate to call myself both his student and his friend.

John Arras and Sahar Akhtar heavily influenced the path my research on this essay took. Classes with both of them in the early stages forced me to broaden my view of the subject. Chapter 5, in particular, would not have been included had it not been for their influence. My approach to the subject of that chapter developed during the Fall 2011 semester, when John arranged for me to teach a course on patents and biotechnology for the UVa Bioethics program. This is only one small example of the generosity John has shown me throughout the years, with both his time and his resources, which is no small thing considering that I was not officially his responsibility. For all this and more, I owe him a huge debt of gratitude.

Many teachers, colleagues, and friends read drafts of the chapters and talked through the issues with me over the three years I wrote this essay. I especially want to thank Mitch Green, Pete Boettke, Virgil Storr, Chris Coyne, Bill Glod, Bill Hasselberger, Sam Duncan, and Jason Craig. Bob Goedert, distinguished lawyer and my good and true friend, took time out of his very busy schedule to read drafts and double-check both my economics and my understanding of intellectual property law. And a special thanks is due to Julia Mahoney, who agreed to read the whole thing and be the external examiner on my dissertation committee on short notice.

My personal debts are far too numerous to mention here, but I want to single out two in particular. Christine Breton put up with the stress-induced bouts of melancholy and frustration writing and research (sometimes) occasioned. It takes a special person to deal with a burrowed-down graduate student in the final stages of their dissertation project, and without her constant love and support I could not have gotten this thing done.

Ten years ago I told my parents, David and Susan, that I wanted to be a philosopher. Many parents would hit the roof, forbid it, threaten disinheritance. My parents simply nodded, asked a few questions, and said “Ok”. Throughout my time in graduate school, they have offered nothing but encouragement and support, and were always there when I needed help, to take a break and come home for a visit, or just to talk on the phone. Their unflagging support was a necessary condition for all of this. One of the happiest moments in my life was calling home to tell them that I got a job and was graduating; the way they reacted to that news reminded me that I had the best parents in the world.

This dissertation is dedicated to all my teachers, and especially those individuals singled out by name in the dedication. Some I’ve lost touch with, one (Mrs. Stewart, my 5th grade social studies teacher) is no longer with us, and some will be in the room when I defend this dissertation. All were essential.

*Brooklyn, NY
June 2013*

Contents

<i>Preface and Acknowledgements</i>	4
1. Against Instrumentalism About Intellectual Property	9
1.1 A Tale of Two Patents	10
1.2 What is Intellectual Property?	16
1.3 Intellectual Property Theory and Policy I: What Philosophers Talk About When They Talk About Intellectual Property	24
1.4 Intellectual Property Theory and Policy II: What Everyone Else Talks About When They Talk About Intellectual Property	30
1.5 A Preview	33
1.6 Some Remarks on Method	37
2. Incentives Arguments for Intellectual Property	41
2.1 The Logic of the Incentives Argument	44
2.2 Two Criticisms of the Incentives Argument	52
2.3 Expanding the Incentives Argument From the “Demand” Side	57
2.4 Investment and Downstream Creative and Innovative Work	62
2.5 The Incentives Argument is Not Just About Incentives	68
3. Labor as the Basis for Intellectual Property Rights	72
3.1 What Does a Labor Theory Have to Show?	74
3.2 Why Labor Matters	80
3.3 The Productive Capacities View	86
3.4 An Improved Labor Theory	90
3.5 But Are Intellectual Property Rights <i>Justified</i> ?	99
4. Intellectual Property or Intellectual Monopoly?	102
4.1 The Monopoly Argument	106
4.2 Two Problems With the Standard Model of Intellectual Property	112
4.3 Digital Sampling, Copyright, and Derivative Works: A Case Study	123
4.4 Reassessing the Monopoly Argument	133
5. Global Justice, TRIPS, and Access to Medicines	136
5.1 The TRIPS Agreement and Access to Medicines	141
5.2 Two Versions of the Access Argument	145
5.3 Global Justice and International Trade Agreements	149
5.4 TRIPS and Global Justice	155
5.5 Reform vs. Aid	159
5.6 Conclusion	162
6. Some Conclusions	164
<i>Bibliography</i>	174

Chapter 1

Against Instrumentalism About Intellectual Property

1. A Tale of Two Patents

In his *First Treatise* on Government, Locke wrote:

Thus far can the busie mind of Man carry him to a Brutality below the level of Beasts, when he quits his reason, which places him almost equal to Angels. Nor can it be otherwise in a Creature, whose thoughts are more than the Sounds, and wider than the Ocean, where fancy and passion must needs run him into strange courses, if reason, which is his only Star and compass, be not that he steers by.¹

To study intellectual property is to enter a world in which it sometimes appears the “busie mind of man” has quit reason, and though it may be an exaggeration to say that the result is a brutality below the level of beasts, it is certainly not hyperbole to say that fancy and passion have run him into strange courses. Consider the following two examples:

US Patent No. 6,004,596: The Peanut Butter and Jelly Sandwich

The JM Smucker company makes and sells a product called the “Uncrustable”, basically a sealed, crustless sandwich with various fillings, which can be stored in the freezer and then microwaved (similar to the “Hot Pocket”). Smucker’s has multiple patents on this item. Smucker’s attempted to patent a version of the sealed crustless sandwich which it described, in its patent application, thusly: “The upper and lower fillings are preferably comprised of peanut butter and the center filling is comprised of at least jelly. The center filling is prevented from radiating outwardly into and through the bread portions by the peanut butter” (Boyle 2010, xi).

In short, Smucker’s filed a patent for the peanut butter and jelly sandwich. As an example of the “prior art” that this version of the sandwich improved upon – such an improvement constituting the all important “novel” or “inventive step”, a necessary condition for receiving a patent – the patent examiner who reviewed the application cited

¹ *1st Treatise of Civil Government* §58 (in Locke 1993).

a cookbook called *50 Great Sandwiches* by Carol Handslip, which describes (among other things) different ways of making peanut butter and jelly sandwiches.

Smucker's was initially granted the patent, and proceeded to issue cease and desist letters to other food companies manufacturing crustless versions of the peanut butter and jelly sandwich. The patent was reviewed, and cancelled. Smucker's appealed, and lost the appeal, thus losing their patent. The absurdity of the situation was not totally lost on the justices hearing the appeal. One of the justices described during oral arguments how his wife often cuts the crust off the bread of his children's peanut butter and jelly sandwiches and squeezes the bread together, so as to minimize the mess his children will make when eating them; he noted that this could be considered a violation of Smuckers' patent.

US Patent No. 4,259,444: Chakrabarty's Oil Eating Bacteria

After finishing his PhD at the University of Calcutta, Ananda Chakrabarty went to the University of Illinois and began studying a group of bacteria known as pseudomonads. Pseudomonads are capable of breaking down and assimilating large, complex organic compounds, such as camphor and elements of crude oil. Chakrabarty identified and isolated the genes that allowed pseudomonads to accomplish this task. After moving on to the research and development unit of General Electric, Chakrabarty became interested in using pseudomonads to break down crude oil into benign proteins (initially the idea was that pseudomonads could turn oil, which was cheap and plentiful in some parts of the world, into edible proteins, which were not). After the sharp rises in oil prices following the 1973 oil embargo and high-profile environmental disasters caused by oil spills, Chakrabarty became interested in using pseudomonads to clean up after oil spills.

The problem was that individual *Pseudomonas* strains only possessed genes that enabled them to break down a small number of the hydrocarbons in crude oil. Chakrabarty figured that a pseudomonad that had all these genes could break down the entire substance, and so he inserted plasmids containing the suite of hydrocarbon-degrading genes into a single strain of *Pseudomonas Putida* and cultivated it in his GE laboratory. The result was a recombinant organism, a genetically modified pseudomonad capable (in theory) of cleaning up oil spills by breaking down crude oil into harmless proteins.

At the time living organisms were not considered patentable, but GE filed for a patent on Chakrabarty's genetically modified pseudomonad, figuring that genetically engineering the organism qualified it as an invention. The patent was initially rejected, but the US Court of Appeals for Patents and Customs overturned the decision. Sydney Diamond, Commissioner of the Patent Office, appealed all the way to the US Supreme Court.

The case of *Diamond v. Chakrabarty* became a landmark. In a close 5-4 vote, the Supreme Court ruled that Chakrabarty's oil-eating bacteria was a patentable invention, as the subject matter of the patent was the modified genome of the organism, and thus was no different from other "composition of matter" patents. For the first time ever, a patent was granted on a living organism. The decision revolutionized the pharmaceutical industry and paved the way for commercial biotechnology, by providing a means for companies to attract investment through the patenting of genetically modified organisms

(or GMOs) and genes. It also created a great deal of controversy, as the idea of owning the “stuff of life” seemed (and still seems) an *infamita* to many.²

These are simply two vivid examples of some of the “strange courses” the evolution of intellectual property has run in recent years. Still considered, at the beginning of the 20th century, state-granted limited entitlements to help protect industries like publishing, intellectual property rights – patents, copyrights, trademarks, and trade secrets – have grown into one of the most important forms of property in the 21st. Intellectual property institutions have a key role in regulating industries as diverse as media, biotech, and environmental design. Even casual observers of world affairs can’t help but encounter stories about the various controversies, developments, and noteworthy events in the world of intellectual property. Consider just some of the stories that have grabbed international headlines in recent years: a giant protest against a bill to stop online copyright infringement; Supreme Court cases about patenting human genes and genetically modified crops; endless and seemingly absurd litigation pitting the world’s two largest cell phone manufacturers against each other; and a strident denunciation and call for action from the President of the United States – issued during a week which also featured massive protests in Turkey, an uptick in violence in the Syrian civil war, Russian naval exercises in the Mediterranean, and controversies over government collection of

² The story of Smuckers’ attempt to patent the peanut butter and jelly sandwich is told in Jaffe and Lerner 2006 and the preface to Boyle 2010. My account of Chakrabarty’s patent comes from Chakrabarty’s autobiographical reflections on the case (Chakrabarty 2002).

cell phone records and internet data – of an obscure practice derogatively known as “patent trolling”.³

Property is a topic of perennial interest to philosophers, but only recently has the discipline begun paying attention to intellectual property. This is in sharp contrast to the level of attention intellectual property receives in disciplines such as economics, legal theory, media studies, public policy, and public health. This, however, is not (wholly) due to the philosopher’s traditional (and largely mythical) inattentiveness to developments in the “real world”. It is very much an open question whether the central issues about intellectual property really have any “philosophical” dimension at all. To be sure, many issues raised by intellectual property intersect with topics of interest to a wide variety of philosophical sub-disciplines and projects, including (but by no means limited to) philosophy of art and aesthetics, political philosophy, social epistemology, applied ethics, philosophy of biology, philosophy of economics, and philosophy of language. Intellectual property raises or touches on many questions pursued by those who work in these areas: What is the difference between a work of art derived from another, and one that is copied from another? What is the nature of authorship, and how much does the act of creation owe to artists’ cultural and historical milieu? Are intellectual property rights even *property*, in the same way rights to land and material objects are property? Are

³ The stories referred to here are: the giant public outcry and protest against the Stop Online Piracy Act (SOPA) in Fall 2011; the cases *Bowman v. Monsanto Co.* and *Association for Molecular Pathology v. Myriad Genetics*, which challenged patents on human genes and GMOs, respectively; the ongoing litigation battle between Apple and Samsung over patents related to their market-leading cell phones (the iPhone and Galaxy, respectively); and a speech given by President Barack Obama, coupled with several White House press releases during the week of June 2-7 2013, denouncing firms that are “non-practicing entities” (that is, firms that hold but do not use patents, and aggressively seek licensing fees from others by threatening litigation based on their portfolio of patents – otherwise known as “patent trolls”).

genetically modified organisms really “inventions”? What is the best ontological account of genes, anyway? What is an incentive, and how much can the concept help to explain and understand behavior?

Though these are certainly interesting issues, they are not the main, overarching question about intellectual property pursued in its different aspects by various disciplines. Simply put, that question is: *should we have intellectual property institutions at all, and if so, what should they be like?* This *uber*-question is almost always framed, within the mainstream of discussions about intellectual property theory and policy, in terms of the *instrumental* value of intellectual property institutions. Specifically, it is framed as turning on whether or not patents and copyrights are necessary to provide incentives for intellectual labor, in order to ensure an adequate supply of literary, artistic, scientific, and biomedical goods. When framed in this way, there doesn't really appear to be much for a philosopher to do. All the relevant parts of this question are empirical.

This reduction of the central issues about intellectual property to questions about the instrumental value of intellectual property institutions – which I will call *instrumentalism* – is the main theme of this dissertation. My goal here is to argue that questions about intellectual property cannot be reduced to questions about the instrumental value of intellectual property rights. There are deep normative issues that we must confront in order to determine not only whether we should have intellectual property, but also what shape intellectual property institutions should take. Instrumentalism is not simply ignorant of this dimension; rather, instrumentalism holds that these sorts of considerations are simply *not relevant* in assessing intellectual property institutions. The relevant variable is the incentivizing role of intellectual property rights,

and the need (or lack thereof) of IP institutions to ensure an adequate *supply* of intellectual labor. Other considerations can be bracketed, and “to have IP or not” can be reduced to a technical question about whether IP is necessary to incentivize, and thus ensure adequate supply of, creative, artistic, and innovative work.

I consider, in discrete but related chapters, the most widely cited and important arguments for and against intellectual property in the professional literature, as well as those floating around in popular and informal discussions. Along the way I touch on a number of topics concerning my central subject. Each of these chapters can be read independently (for the most part) of the others, but they are nevertheless related. All explore aspects of the genuinely *philosophical* dimension of intellectual property theory and policy, and aim to show that an answer to that question involves serious philosophical reflection on topics from the nature of creativity to how demands of global justice shape international trade agreements.

2. What is Intellectual Property?

The term “intellectual property” (or “IP”) is a blanket term for a number of legal entitlements, which are collectively referred to as “intellectual property rights” (or “IPRs”).⁴ There are four entitlements in particular that are paradigmatic:

- A **copyright** gives its holder exclusive rights to make and distribute copies of creative or artistic works, including (but not limited to) novels, poems, plays, architectural designs, sounds, images, software, musical compositions, movies,

⁴ For the remainder of the dissertation, I will abbreviate “intellectual property” as “IP” and “intellectual property rights” as “IPRs”. What the abbreviations lack in elegance they make up for in brevity.

sculptures, and radio and television broadcasts. In the United States, the term of a copyright is currently the duration of the life of the author plus seventy years after death. Of crucial importance in understanding copyright is the distinction between the protected idea, and *expressions* of the idea. Copyright gives its holders rights over *expressions* of ideas, but not ideas themselves. The best illustration of this is the famous case of *Baker v. Selden*.⁵ Charles Selden, in 1859, copyrighted a book called *Selden's Condensed Ledger, or Book-keeping Simplified*. The book laid out a simplified accounting system that Selden had devised. After his death, another book, by W.C.M. Baker appeared, which presented what was, in effect, Selden's system, but in a different and more extensive manual. Selden's widow sued for copyright infringement. The court held that Baker had not committed a copyright violation, because copyright protected not Selden's idea (the accounting system), but rather the manual Selden had written. In other words, not the idea, but the expression of the idea in the how-to manual, was covered by Selden's claim.⁶

- A **patent** gives holders exclusive rights over a variety of inventions, such as machinery, pharmaceutical formulas, new varieties of seeds, devices, and processes. Unlike copyright, patents must be applied for (any copyrightable idea, as long as its fixed in a medium, is copyrighted, no application necessary). In order to get a patent, an applicant has to describe, in detail, the invention, and

⁵ 101 US 99 (1879). The case was a landmark in US intellectual property law history. For discussion of the case, and the so-called “idea/expression” dichotomy generally, see Litman 1990 and Yen 1989.

⁶ The court did however hold that a *patent* could potentially give one exclusive rights to the “useful arts” described in such a manual.

show that it meets certain criteria. Most importantly, the invention must be “novel” and “nonobvious”; that is, it must be a genuinely new invention, and cannot be “obvious” to anyone who has the relevant technical knowledge.

Patents last for a period of twenty years after their filing date. This means that the effective life of a patent – the amount of time patented inventions are on the market – is sometimes less than the actual life of the patent.

- **Trade secret** protection, unlike patent, does not require public disclosure of information (hence the name). Trade secret laws vary greatly by jurisdiction, but in general, they give companies means to keep economically advantageous information secret, such as requiring employees to sign non-disclosure agreements or protections against industrial espionage. The classic example of a trade secret is the secret formula for Coca-Cola. However, if the information is acquired lawfully, such as through reverse engineering, trade secret laws no longer apply.
- **Trademark** protects brand names, company logos, or phrases closely associated with a company or product. Trademarks allow companies to protect their brands from counterfeiting, and maintain the signaling function of brand designators such as logos (like the now ubiquitous Nike “swoosh” symbol).⁷

⁷ Other rights, such as rights over use of one’s public likeness or persona, so-called “information rights” such as rights to keep certain information (like medical records) private, and the misleadingly titled “moral rights” of creators of artistic and literary works to, among other things, attribution as creator of a work are also sometimes grouped with patents, copyrights, trademarks, and trade secrets as IPRs. I won’t touch on these here, but will restrict myself to the paradigmatic quartet.

The distinguishing feature of IPRs is their subject matter. The subjects of patent, copyright, trademark, and trade secret are various abstracta: signs, images, designs, formulas, inventions, novels, software code, and the like. Though the term is far from ideal, throughout the dissertation I will refer to the subject matter of IPRs as *ideas*. It is very important, however, to remember that even though the subjects of IPRs are ideas, IPRs do not give their holders exclusive rights *to ideas*, but exclusive rights to certain *uses* of ideas. A copyright, for instance, gives its holder rights to make and distribute copies; it does not give its holder rights to control what owners do with those copies after they buy them (the technical way to phrase this is that a copyright holder's rights are exhausted at first sale; patent rights are not exhausted at first sale). What uses are included in the package of rights granted by different IPRs varies, and there are serious questions we can ask about what rights *should* be included, such as whether copyright should prohibit digital sampling of recorded music. We will revisit this distinction, between rights to ideas and rights to uses of ideas, in Chapter 4.

The vagaries of the particular entitlements referred to as “IPRs” are not really going to concern us here. Most of the issues discussed will stem not from the particulars of patent or copyright, but from the general features of all IPRs (the exception is Chapter 5, which will focus exclusively on pharmaceutical patents, and some of the issues raised by the particulars of international trade laws governing patented pharmaceutical products). So, for the remainder of the dissertation, we can distinguish between, on the one hand, actual, existing IPRs, such as patent and copyright, and the *core idea* of an IPR. The core idea of an IPR is a set of *exclusive, transferable* rights to certain *uses* of ideas.

We need to distinguish, as well, between IPRs, on the one hand, and the broader category of IP *institutions* on the other. I mean “institution” here in the sense defined by Rawls, as “...a public system of rules which defines offices and positions with their rights and duties, powers and immunities, and the like. These rules specify certain forms of action as permissible, others as forbidden; and they provide for certain penalties and defenses, and so on, when violations occur” (Rawls 1971, 55). In the broader category of IP institutions we have to include not only IPRs, but also things such as national patent offices and patent application and review processes, courts which consider issues of IP law and prescribe remedies or other punishments, offices such as patent examiner, and trade agreements governing IP, among other things. Sometimes the issues discussed here specifically are about IPRs, and sometimes they are about this broader category; Chapter 5, for instance, specifically considers an international trade agreement, part of the set of global IP institutions.

IP has a long history, dating back to antiquity.⁸ The ancient Greek colony of Sybaris granted IPRs on culinary recipes and other luxury items. The Republic of Venice enacted the first patent laws in the sense defined above in 1474, which allowed for state enforcement of exclusive uses of inventions as long as the ideas were publicly disclosed. England’s first patent law was the 1624 Statute of Monopolies, which required inventions to be genuinely novel as a condition for grant of a crown monopoly on their manufacture, sale, and use. England’s 1710 Statute of Anne is generally considered the first copyright

⁸ For this brief history, I’ve drawn on the following sources: Boyle 2010; Johns 2011; Litman 1990; Machlup and Penrose 1950; Mossoff 2007; Raustiala and Sprigman 2012; and Rose 1995.

law ever, which gave authors exclusive rights to license their works to a printer, and printers exclusive rights to make copies of works for a period of fourteen years.

Prior to ratifying the federal constitution, different US states passed their own patent laws, largely modeled on the British. The US constitution made IP law a federal matter, by explicitly giving Congress the power to grant IPRs – in the now famous and ubiquitous phrase – “for the promotion of science and the useful arts”. Various European countries enacted their own patent and copyright laws throughout the late 18th and 19th centuries. Of particular note is the creation of the French patent system during the French Revolution in 1791, which originally did not require patent applications because inventors were considered to have natural rights over their inventions. IP laws outside of Europe and the United States usually followed as a result of European colonization, and took on various forms and idiosyncrasies. Indian patent law, for instance, until recently did not allow patents on pharmaceutical products, only processes for making medicines.

In the 19th century, the economic, political, and cultural importance of IPRs began to grow exponentially. The 19th century brought the first attempts to standardize international IP laws. The Berne Convention of 1886, and the Paris Convention for the Protection of Industrial Property of 1883, established trade rules for copyrighted and patented goods, respectively. These were the forerunners to the landmark Agreement on Trade-Related Aspects of Intellectual Property Rights, or TRIPS agreement, adopted during the negotiations that led to the creation of the World Trade Organization in the 1990’s (we will discuss this more extensively in Chapter 5).

It is customary to view the explosion of importance of IPRs, as well as phenomena such as the TRIPS agreement, as distinctively late 20th century developments,

but its important to see that these are a continuation of trends that began in the 19th century.⁹ As world economic development took off following the industrial revolution, the importance of IPRs grew apace. I'm not going to speculate about it here, but it is worth considering whether there is an interesting correlation between the growth in importance of IPRs and the diversification of economies and markets that comes with development. IPRs continued to become more and more economically significant as new industries, such as the motion picture industry, entered the scene throughout the beginning of the 20th century, and there is no doubt their present importance is due in large part to the status of industries such as digital media, information technology, and pharmaceuticals in the world today. One of the more interesting arguments for IPRs (touched on in Chapter 3) is that they help facilitate a division of labor. By providing a specific way for artists and inventors to make a living, and certain kinds of firms a way to raise capital, IPRs allow these firms and individuals to specialize in, and devote themselves full time to, endeavors like medical research. I don't mean to suggest here that IPRs are a *cause* of this development, only that there does appear to be a connection between the internal diversification of economies and the existence of institutions that allow for the specialization that helps facilitate this diversification, one that might be an interesting topic for economic historians.

The present status and importance of IPRs is shaped by the general trend of interpreting these rights broadly and favorably. Public opinion has swung in different directions about IP. The depression of 1890 in the United States, for instance, resulted in a tide of anti-IP opinion, as patents especially were seen as just another form of

⁹ For examples the claim that the expansion of IP is a distinctively late 20th century phenomenon, see Boyle 2003 and Lemley 2005.

monopoly. This was repeated during the Great Depression, when courts showed a penchant for invalidating patents on appeal. The United States has – at least until very recently – been riding a high tide of pro-IP opinion. The last 40 years has seen an unprecedented expansion in interpretation of both the scope and strength of IPRs, as well as the subject matter protectable by patents and copyrights. Things previously thought off limits, such as business methods (like Amazon’s “1-click” e-commerce technique), software, GMOs, discrete elements of musical recordings (like single snare drum cracks), and human genes are now considered patentable or copyrightable material.¹⁰ At the same time, more and more ideas come under some form of IP protection. James Boyle has called this twofold expansion of IP, derogatively, the “second enclosure movement” (Boyle 2003), and whether we agree with his negative opinion or not, there’s no doubt it is a significant development in IP history.

Whether the high tide will continue or ebb is debatable. The twofold expansion of IP has definitely yielded some ugly moments, such as the lawsuits filed against teenagers by major recording companies over their illegal downloading of digital music in the early 2000’s. There are signs that public opinion is shifting decidedly against IP. The massive protests that derailed the Stop Online Piracy Act in Fall 2011, for example, as well as the huge grassroots movements in Europe against biotech may well signal the end of the pro-IP tide of the last decades.¹¹

¹⁰ The title of an article by Robert Merges on this expansion of subject matter – “As Many as Six Impossible Patents Before Breakfast” (Merges 1999b) – perfectly captures the mood generated by the phenomenon.

¹¹ As I was writing this chapter, the US Supreme Court issued its decision in the case of *Association for Molecular Pathology v. Myriad Genetics*; the court ruled 9-0 against Myriad Genetics’ claims regarding patents they hold on the BRCA1 and BRCA2 genes,

3. Intellectual Property Theory and Policy I: What Philosophers Talk About When They Talk About Intellectual Property

Almost from the very beginning, the central theoretical and policy questions about IP were framed as questions about the tradeoffs involved with granting IPRs. IPRs in Britain and the US were conceived as a kind of state created and enforced monopoly, akin to things such as the British East India Company, and discussions about IP took their cues from broader discussions in about monopoly. The dynamic was: monopolies are bad, but IPRs might well be a necessary evil if the tradeoffs (the good effects of IPRs, for the bad consequences of giving anyone a monopoly) are worth it.

This theoretical framework dominated discussion about IP throughout the 19th century. In a famous and oft-quoted letter, Thomas Jefferson perfectly laid out the ambivalence of his age over IP:

If nature has made any one thing less susceptible than all others of exclusive property, it is the action of the thinking power called an idea, which an individual may exclusively possess as he keeps it to himself; but the moment it is divulged, it forces itself into the possession of everyone, and the receiver cannot dispossess himself of it. Its peculiar character, too, is that no one possess the less, because everyone possess the whole of it; as he who lights his taper at mine, receives light without darkening me. That ideas should freely spread from one to another over the globe, for the moral and mutual instruction of man, and improvement of his condition, seems to have been peculiarly and benevolently designed by nature. (Boyle 2010, 20)

One would be hard pressed to find a more eloquent anti-IP statement. Of course,

Jefferson continues:

Society may give an exclusive right to the profits arising from [inventions], as an encouragement to men to pursue ideas which may produce utility, but this may or

human genes believed to be instrumental in the development of certain breast and ovarian cancers. The decision looks like a lock to be a landmark in the history of patent law. The upshot of that decision is that, with some important technical exceptions, human genes are no longer considered patentable subject matter, further evidence of a (possible) rollback in the recent tide of pro-IP sentiment. See Adam Liptak, “Justices, 9-0, Bar Patenting Human Genes”, *New York Times* Jun 14, 2013.

may not be done, according to the will and convenience of the society, and without complaint from any body...other nations [besides England] have thought that these monopolies produce more embarrassment than advantage to society; and it may be observed that the nations which refuse monopolies of invention, are as fruitful as England in new and useful devices. (Boyle, op. cit.)

In this letter, Jefferson perfectly lays out the theoretical framework that continues to dominate discussion about IP to this day (as for the equating of IP with monopoly, we will return to that vexed issue in depth in Chapter 4). This framework has two elements:

- There are two independent and sufficient lines of justification for IP. Either IPRs are a form of property right, and arguments for property rights can justify IPRs, or they are a means to generate incentives, and are instrumentally justifiable. If they are not the former (as Jefferson thinks) they may still be the latter (or neither, as he implies at the end), but each line of justification is independently sufficient.
- The chief rationale for IPRs – assuming no arguments for property rights can be extended to cover patents and copyrights – is the “encouragement”, or incentives, they provide in order to get individuals to engage in “utility” producing intellectual labor. If they fail either to provide the right incentives, or ensure an adequate supply of intellectual labor, then they are not justifiable.

The combination of these points makes up the standard view of IP. By “the standard view” I don’t mean that all IP theorists are, like Jefferson, skeptics who think that IP is, at best, a necessary evil (though many are). Rather, what I mean is that most IP theorists see the question in roughly the same way that Jefferson does. Debate about IP centers around two central questions posed: are IPRs a form of property right, justifiable along similar (moral) grounds as other property rights? And if not, are they instrumentally

justifiable, as means to provide incentives for (presumably socially beneficial) intellectual labor?

Most philosophical discussion of IP is concerned with the first question.¹² “Yes” answers to the first question take the form of extensions of classic theories of property to cover the subject matter of IPRs. The most widely discussed are *Lockean* theories of IP.¹³ The central idea is that the same Lockean considerations that apply to property rights in physical objects also apply to property rights in ideas. Artists, writers, and inventors “mix their labor” with the images, novels, and machines they create; the same considerations that apply to mixing one’s labor with physical things should also apply (assuming it does at all), *mutatis mutandis*, to ideas as well. Some defenders of Lockean theories of IP consider ideas especially amenable to a Lockean justification, generally for two reasons. First, because ideas are *nonrival* appropriation from the intellectual commons satisfies the famous Lockean proviso to leave “enough and as good” in the commons better than appropriation of material objects.¹⁴

The property of nonrivalrousness is one of the most important concepts in the theory of IP, and it will come up in just about every single chapter in this dissertation. We will discuss nonrivalrousness in depth in Chapters 2 and 4, but for now, the basic idea is: a good is *nonrival* if there is no or at least a negligible cost generated by using and

¹² There are notable exceptions, the most significant of which is Thomas Pogge’s work on pharmaceutical patents (Pogge 2007; 2008, Ch. 9).

¹³ Lockean theories of IP are developed and presented by: Child 1997; Gordon 1993; Moore 1997; and Yen 1990. For discussion and criticism of Lockean views, see: Becker 1993; Hettinger 1989; Hughes 1988; and Shiffrin 2001. Locke himself wrote a short essay on literary property and publishing; see Locke 1997 (but see also Shiffrin 2001, for why we shouldn’t read too much into Locke’s remarks).

¹⁴ See Hughes 1988 and Moore 1997 especially for discussion of this point.

accessing it. Using some good can generate costs in many ways: it can reduce the value of the good by depleting it (for instance, removing timber from a forest), it can preclude others from using the good (when I drive my car, no one else can drive it at the same time), or it can cause congestion and thus reduce the utility of using the good (for instance, the more drivers there are on a popular road, the less convenient a route that road becomes).

Ideas are paradigmatic examples of nonrival goods. Adding additional users to a good does not generate any additional costs. It does not reduce, deplete, crowd out, or preclude anyone from reading *Moby-Dick* if 100 or 1000 people are also reading the book at the same time I am. So, when I “appropriate” an idea from the intellectual commons – say, when I improve on the internal combustion engine and make a more efficient version of it – I do not preclude others from also tinkering with the engine, or from using the same ideas for a variety of other tasks. When my patent on my new engine expires, others can make and market versions of my engine without any overall reduction in the value of the ideas. Appropriating elements of the design for the internal combustion engine to make a new invention seems to leave not just “enough and as good” in the intellectual commons; it seems to leave the intellectual commons largely unchanged with regard to its value.

The second reason why IP is considered (by some) so amenable to a Lockean theory of property is the seemingly tight connection between the individual and the product of their labor. When a writer draws on her experiences growing up in the Lubavitcher community of Hasidic Jews in Crown Heights, Brooklyn to write a play, the resulting work contains even more elements of herself, and is (again, seemingly) even

more intimately connected with her mind and labor than physical objects produced through manual labor.¹⁵ The play is, in the words of Alfred Yen, “practically an extension of the author herself” (Yen 1990, 547), in a tighter and more intimate way than, say, a chair she built out of lumber.

This alleged tight connection between creators and inventors is the impetus for another “yes” answer to our first question, *personality-based* theories of IP.¹⁶ The central idea of such theories is that inventive or creative works are an extension of the author or inventor’s personality. These works embody or realize elements of her personality and mind, and express parts of her perspective on the world, and so lack of control over them is alienating. IPRs are ways of ensuring the requisite level of control, by placing others under duties to refrain from interfering with the author or inventor’s control over the products of their intellectual labor.

Related to personality views, but offering a different set of moral foundations altogether, are *Kantian* theories of IP.¹⁷ These are “Kantian”, in that they take their cues from Kant’s views on public reason (primarily in *Metaphysics of Morals*, “What is

¹⁵ There is a great deal of justifiable skepticism about the notion of authorship behind this claim. See Hettinger 1989 and Litman 1990 for criticism.

¹⁶ Sometimes these views are infelicitously called “Hegelian” theories of IP, as Hegel is cited as their intellectual godfather, on the basis of remarks scattered throughout ¶¶59-70 of *Elements of the Philosophy of Right* (Hegel 1991). Schroeder 2006 presents a “Hegelian” theory of IP. Personality-based views of IP are associated closely with the work of Margaret Jane Radin on property: see especially Radin 1982. For discussion of personality theories, see Hughes 1988.

¹⁷ Laura Biron and Anne Barron have both developed Kantian theories of copyright (Barron 2012; Biron 2012). The main source for these views (besides Kant himself) is Onora O’Neill’s work on Kant and public reason (especially O’Neill 1990, Chs. 1-2).

Enlightenment?”), and two short essays he wrote about publishing).¹⁸ Like personality-based theories, Kantian theories of IP begin from the close connection between the products of intellectual labor and the author’s or inventor’s mind and personality. However, for Kantian theorists, the role of IP institutions is not to ensure that authors and inventors have control over the products of their intellectual labor (or at least, not just that), but to structure expression so that it is in accord with norms of communicative rationality and public reason. IP institutions are not “property”, but rather are institutions that realize communicative norms, and so ensure a robust public sphere structured in accord with these norms. This may involve property-esque entitlements over works (such as some kind of copyright), but the ultimate moral foundations for these are not a theory of property but the Kantian theory of public reason.

I will discuss Lockean theories briefly in Chapter 3, where I present a reconstructed version of a labor argument for IP I call the productive capacities view.¹⁹ I do not discuss either personality or Kantian theories of IP at all in this essay, however. My goal is to address questions discussed by those working on IP theory and policy generally, not just convince other philosophers that a Lockean or Kantian theory of IP is correct and other accounts are not. To many philosophers, I understand this will make this essay any or all of the following: (a) misguided, (b) un-rigorous, (c) a betrayal or dodge, as it neglects the “central philosophical issues” involving IP, (d) boring, and/or (e) “applied” (as opposed to “real”) philosophy. *C’est la vie* (or as we say where I grew up

¹⁸ These are “On the Wrongfulness of Unauthorized Publication of Books” and “On Turning Out Books” (both in Kant 1999).

¹⁹ However, as I point out there, I do not consider my view – and many Lockean would not consider my view – a Lockean theory of IP, despite similarities with such views.

in Chicago, “it is what it is”). I will say something about my approach to the problem in the final section of this chapter; but I hope overall the reader will let the results speak for themselves, and leave off judgment until the whole essay is on the table.

4. Intellectual Property Theory and Policy II: What Everyone Else Talks About When They Talk About Intellectual Property

It is generally assumed that if we reject any distinctly “moral” foundations for IP, in the form of a Lockean, personality-based, or Kantian justification of IP, then the question turns entirely on the instrumental value of IPRs.²⁰ If IPRs are justified, it is because they provide incentives for the creation of goods we desire, and serve a useful role as institutions for regulating technology and media. The vast majority of work in IP theory and policy outside of philosophy (and some within) proceeds as if the answer to the first of the two questions discussed above is “no”, without much (or any) discussion of the sorts of considerations in the previous section.²¹

We will discuss the primary instrumental argument for IP – the incentives argument – in depth in the next chapter, and return to elements of it in Chapter 4. In brief, the argument is: producing ideas is very costly. It takes time, energy, resources, and a

²⁰ Sometimes the term “utilitarian” is used in the legal literature on IP, instead of “instrumental”. This is very unfortunate, as IP “utilitarianism” (*sic*) is usually meant to be an alternative to “moral” justifications of IP (whereas an *actual* utilitarian justification of IP would not be an alternative, but merely another version of such a moral justification). I have used the terms “instrumental” and “instrumentalism” throughout, to avoid this confusion. The reader should note, however, that they could encounter the term “utilitarianism” in the legal literature on IP where I would use the term “instrumentalism”. There is a great deal of work in law and economics that could be considered a utilitarian (no square quotes) theory of IP; the work of William Landes and Richard Posner is a good candidate (Landes and Posner 2003).

²¹ A major exception to this is a handful of legal theorists such as Robert Merges, Wendy Gordon, and Alfred Yen (see Gordon 1993; Merges 2011; Yen 1990).

great deal of skill and talent to write a new novel, design a building, or an innovative automobile brake system. But ideas have two properties that make it difficult for intellectual laborers to get any return on their investment through unprotected release of the products of their labor into the market. The first property is nonrivalrousness. Since ideas are nonrival, copies or tokens of ideas – engines that realize a design, or units of a cancer drug, or copies of a novel – can be produced at will, without depleting or subtracting from the value of the design, formula, or novel. This in and of itself isn't a problem. Ideas, however, have a second property – they are not only nonrival, they are also *nonexcludable*. This means that, once an idea has been made available to one person, it is very difficult to prevent it from being available to anyone. If I write a novel and sell a copy of it to a friend, for example, its difficult for me to prevent that person from making copies of the novel and selling it for less than my price to others who may be interested in reading it.

For reasons we will examine in detail at the beginning of the next chapter, because ideas are nonrival and nonexcludable, it is (allegedly) very difficult to get everyone who benefits from them to internalize a share of the cost of their production through sale of novels, drugs, and engines (sans IPRs). IPRs provide a solution to this problem. IPRs prohibit everyone who is not their holder from making and marketing copies or tokens of the ideas that are the subject of the IPR, and allow IPR holders to use the law to enforce these rights and seek damages. The effect is to give IPR holders the chance to sell their goods at a price higher than marginal cost for the effective life of the IPR, and this chance at profit (so the argument goes) incentivizes intellectual labor and, more broadly, investment in the production of drugs, engines, novels, and the like.

Vast swaths of Canadian forest have been leveled to make the paper on which all the different discussions, criticisms, defenses, epicycles, and versions of the incentives argument are printed. Anyone even casually familiar with the literature on IP – in disciplines *other* than philosophy – has likely encountered rehearsals of the argument many times. We will discuss criticisms of the incentives argument in Chapters 2 and 4. For now I want to point out that, though the incentives argument is supposed to be instrumental, and is supposed to reduce questions about IP policy to questions about the incentivizing function of IPRs, it is in fact awash in implicit normative claims and commitments. Foremost among these assumptions is that we all, in general, have a compelling interest in having a certain supply of intellectual labor and investment in intellectual goods, and this interest is sufficiently robust to commit public funds and resources to enforcing property rights if, in the absence of these rights, that supply would be sub-optimal. This is far from obvious. It's certainly not clear that the interest many people have in blockbuster action movies, romance novels, and pop records justifies the burdens on many communities of artists of copyright (on one example of this, see the discussion of digital sampling in Chapter 4), let alone the public expense of adjudicating copyright claims. Even if that expense is negligible compared to other budget outlays, it is still an expense, and it is reasonable to say that those who would gladly trade less pop music for more public funding of health care for the poor are owed a justification.²²

²² This is one place where instrumentalism differs from any potential utilitarian justification of IP institutions (and one reason why “utilitarian” is not a synonym for “instrumentalist”). A utilitarian justification would presumably have reasons why uses of public resources to support IP institutions give a bigger net boost to utility than greater spending on, say, health care. A utilitarian theory would make no normative *assumptions* about the value of such a tradeoff; rather it would entail a normative *commitment* to prioritizing supply of good ideas over something like increased spending on health care,

5. A Preview

The normative assumptions implicit in the incentives argument, and the way they affect IP theory and policy, is the starting point for this essay. In the next chapter, I go in depth into the incentives argument, and focus especially on two recent and very sophisticated criticisms of it. These are interesting not just because of their sophistication, but because they are representative of broad sentiments about IP expressed and discussed in both professional and popular debates. The central claim of my discussion in the next chapter is that we have to assess IP institutions not just from the “supply” side – that is, in terms of whether they help produce the optimal supply of intellectual labor and investment – but also from the “demand” side – in terms of whether they also incentivize production of the right kinds of goods. Determining what counts as the “right kinds” of goods, however, means going into a host of normative issues: What kinds of targets should medical research aim at? What artistic and creative goods should we incentivize production of and investment in? How do we determine whose preferences we treat as relevant and whose we discount in determining what kinds of goods IP institutions should incentivize? Or more broadly: what kind of art, science, culture, technology, and medicine *should we have*, and what is the role – if any – of legal institutions like IP in shaping and regulating these areas of endeavor?

if the former increases utility more than the latter. In fact, utilitarian views wouldn't even raise these sorts of issues – about favoring certain interests or preferences over others – in the first place. There would be no question about the way in which IP institutions favored certain interests of certain individuals over others, because on a utilitarian view these sorts of considerations are morally inert; what matters is whether IP institutions increase overall utility, not who would be the primary beneficiaries of that increase – another example of how utilitarianism does not respect or account for the “separateness of persons” (Rawls 1971).

It's very important to note here that the goal of this argument is not *just* to point out that there are normative issues involved in assessing IP institutions; that is, after all, a fairly pedestrian point. The point of the argument is also to show *where those considerations enter in*. The central normative question, according to Chapter 2, is not whether the tradeoffs between the costs and benefits of IP are worth it; rather, it is what counts as success or proper functioning when we are looking at whether IPRs generate sufficient incentives in the first place.

In sum, assessing IP institutions from the “demand” side is an irreducibly, and inescapable, normative problem, and so we cannot – even if we reject “moral” arguments for IPRs – reduce IP theory and policy to a technical engineering question about the incentives structure created by IPRs. But of course we shouldn't eschew “moral” arguments for IPRs. In fact, the way the dichotomy between moral and instrumental arguments is set up should also be rejected wholesale. In Chapter 3, I consider the most popular “moral” argument for IP: Lockean-style labor arguments.²³ The goal there is not to ask whether a labor argument is sufficient justification for IPRs. Rather, the question posed there is whether or not a version of the labor argument gives reasons to accord a certain weight to the interests of intellectual laborers. I present a reconstructed version of the labor argument, which I call the *productive capacities view*, in support of the claim that the interests of intellectual laborers in controlling the product of their labor should carry weight in determining whether we should have IP institutions.

As presented in Chapters 2 and 3, the incentives and labor arguments are not opposed and independently sufficient justifications of IP, but are complementary. Part of

²³ By “popular”, I mean popular in technical literature and public discussion of IP, not just the philosophical literature.

figuring out what incentives structure IP institutions should create, and thus what shape IP institutions should have, involves figuring out what interests of which individuals are relevant. The reconstructed labor argument in Chapter 3 gives a reason to give a certain weight to the interests of intellectual laborers. This, as I point out in the last section of Chapter 3, is *not* a sufficient justification for IPRs. The argument in Chapter 3 gives only a positive case for IPRs; no argument is given as to how to rank or prioritize the interests of intellectual laborers with regard to the interests of other individuals in having looser or no IP institutions.

In Chapters 4 and 5, I turn to two very strong arguments against IP, in an effort to see whether or not there are feasible IP institutions that can either mitigate or remove the harms critics lay at IP's door. In Chapter 4, I examine the *monopoly argument*, which holds that IPRs are equivalent to a monopoly over markets for goods that realize protected ideas, and come with all the harms, costs, and general shenanigans associated with monopolies. The monopoly argument is almost as old as IP theory itself, stemming from debates in late 18th and early 19th century political economy about free trade and mercantilism.

I argue in Chapter 4 that IPRs are not monopolistic by nature. The key to the monopoly argument is what the economists Michele Boldrin and Danny Levine call *downstream licensing privileges*: the power IPRs give their holders to control what people who buy tokens of ideas (such as individual copies of novels) do with those ideas and goods after sale (Boldrin and Levine 2008). The central question in assessing the monopoly argument is thus whether the two components of IP are necessarily, or just accidentally, bundled. That is, whether or not the package of exclusive rights to uses of

ideas and the downstream licensing privileges that would make IPRs monopolies come apart. I argue in Chapter 4 that IP is not monopolistic by nature. Downstream licensing privileges are not an essential component of IP.

In order to reach this conclusion, we have to go deeper into the incentives argument, and the assumptions about property and value that are behind it. These further meditations on the incentives argument are as important as the conclusions about monopoly and IP reached in Chapter 4. Both the pro-IP incentives argument and the anti-IP monopoly argument are based on a twofold confusion about the nature of IP. I argue in Chapter 4 that, when we remove this confusion, we see that IPRs are not, like old-fashioned property rights in land and material objects, rights to control use and access to ideas. Rather, they are rights to *extract units of value* from *rival, renewable, and nonexcludable goods*, and so are more akin to rights to cut timber from a forest or remove swordfish from a fishery than rights to land. This has serious consequences for assessing both the monopoly and incentives arguments, and for IP theory generally.

The heart of Chapter 4 is a small case study of digital sampling of recorded music and copyright, as a demonstration of the main point. Chapter 5 is an extended case study, of pharmaceutical patents and the effects of the Agreement on Trade Related Aspects of Intellectual Property Rights (or TRIPS agreement) on access to medicines around the world. I argue there for two points. First, criticisms of TRIPS, on the grounds that the agreement's distributive effects show it fails to meet requirements of global justice, are right, but for the wrong reasons. TRIPS fails to meet minimal standards of global justice not (just) because of its distributive effects, but because the internal structure of the agreement shapes global pharmaceutical markets in a way that is fundamentally unjust.

Second, as a consequence of this first point, it's not enough to augment TRIPS with aid projects to correct these distributive effects. TRIPS itself must be substantially reformed, or thrown out altogether. In essence, Chapter 5 is concerned with a particular downstream licensing privilege extended to holders of pharmaceutical patents by TRIPS – the power to control manufacturing and sale of generic versions of drugs in foreign countries. Chapter 5 argues that extending this downstream licensing privilege is unjust. Reform of the TRIPS agreement is necessary to trim this particular privilege from pharmaceutical patents, and bring those rights in line with principles of global justice.

6. Some Remarks on Method and Orientation

I began working on this dissertation in Spring 2010. Since then, I've had many conversations where the interlocutors expressed puzzlement as to how this essay counted as "philosophy". Some of these interlocutors were professional academic philosophers, including some who work in my own department – which is dismaying, for a student in the midst of their dissertation project, to say the least. This isn't, however, a facetious question. The subject of the essay is very different from most of the topics philosophers spend their time on, and even amongst the small minority who do care about IP, my focus in this essay is different. It's perfectly legitimate to ask: in what way are these chapters *philosophical* reflections on IP?

Most professional academic philosophers would consider applied ethics the natural home for a dissertation like this. The answer to the above meta-question would then be: "This is a philosophy dissertation, because it's an *applied ethics* dissertation". However, I do not think this essay is, *methodologically* speaking, really an exercise in

applied ethics, as the majority of philosophers understand that subject.²⁴ Though many parts of the dissertation are concerned with ethics, ethics is far from the only – or even the central – topic of the dissertation. The two central chapters, 2 and 4, are concerned mostly with topics we’d consider part of the philosophy of economics and the philosophy of law, and a big part of Chapter 4 – the case study of sampling and copyright – deals with a question in the philosophy of art and aesthetics (namely, the relationship between derivative art and the works it derives from). Fully half the dissertation, then, has little to do with “ethics” (or at least the understanding of the subject embodied in the term “applied ethics”).

The idea of “applied ethics”, as usually understood in philosophy, implies a certain division of labor. Namely, some philosophers figure out what makes actions right and wrong, and some “apply” that theory to particular situations in the “real world” where people have to figure out what they should and shouldn’t do. In that sense of “applied”, this dissertation is not an exercise in applied philosophy, either. There is no theory here, of property or anything else, which is “applied” to the particular “real world” case of IP. It’s more accurate to see this essay not as an application of any particular philosophical theories or views, but rather as the application of a set of methods or intellectual tools characteristic of philosophy to IP.

²⁴ For what it’s worth, I greatly prefer the term “practical ethics” to “applied ethics”, as the sort of business usually called “applied ethics” doesn’t always (or even usually) have anything to do with “applying” previously worked out normative ethical theory. The goal of practical ethics is to sort through normative issues that arise *in the practice of* things like medicine, business, agriculture, and public policy. This requires more than just figuring out what a “Kantian” or a “utilitarian” would say about a particular real world ethical dilemma, which is what is implied by the term “applied ethics”. Many parts of this dissertation are certainly an exercise in practical ethics, as I understand it. Thanks to John Arras, for pushing me to clarify this point.

I believe this essay is methodologically more akin to those parts of the philosophy of science that deal with particular sciences, such as philosophy of physics or philosophy of biology, rather than applied ethics. Of course little to nothing in this essay has to do with the *subject matter* of the philosophy of science. What makes it more akin to work in philosophy of physics or biology rather than applied ethics is a similarity of *approach*. Philosophers of physics, biology, or psychology deal directly with problems that are faced by real working scientists attempting to advance their respective fields. The goal of this work is to help advance science, by helping to work through those conceptual and normative knots that scientists confront. As such, something like the philosophy of physics or biology accepts constraints on what counts as a “good” philosophical theory, in order to ensure that work in philosophy of physics or biology is useful to actual physicists and biologists. These constraints are different from those you would accept if you were trying to develop a philosophical theory of science in general. Thus, there is a sharp distinction between, say, the way in which time is treated in philosophy of physics and speculative metaphysics. The former will rule out conceptions of time that are intuitive but inconsistent with best practice in physics, whereas the latter will not (or at least, will not *necessarily* rule these out).²⁵

My approach here is the same. The goal is not to develop philosophical views about IP for its own sake, but to deal with some philosophical issues raised by IP, in an attempt to contribute to the broader multi-disciplinary discussion about IP theory and policy. The essay will be successful if it can help work through those issues. This is,

²⁵ My views on method in philosophy of science – and thus in this essay – are due to the influence of William Wimsatt and his work in the philosophy of biology (see especially Wimsatt 2007; for discussion of Wimsatt’s project in general, see Griesemer 2011).

again, not to say that the merits (or lack thereof) of Lockean or Kantian theories of IP are not an interesting topic. But I've elected to focus on something else here.

No one discipline owns the theory of IP. Necessarily, the following essay will engage with work in other disciplines, primarily in economics and legal theory. I am not an expert in either of these, and even after ten years of graduate study I barely qualify as competent in philosophy. For what its worth, I think its nevertheless possible to make a contribution to *common* problems, at the intersection of all these different disciplines, even if one is not an expert in any particular one (or, rather, can operate at a high level in *only* one). This essay is not about the economics, philosophy, or legal theory of IP, but rather about a set of issues that are at the intersection of all of these (and more), and not reducible to any. This is not a new thing under the sun. Once there was a discipline called *political economy*, which encompassed parts of what we now consider philosophy, economics, law, sociology, political science, and public policy, and dealt with problems common to them all. This, then, is an essay in the political economy of IP, with a focus on those aspects of its central questions that are most heavily philosophical. The purpose is not to develop a philosophical theory of IP for its own sake, but to advance the political economy of IP, and so contribute to the conversation about IP – and thus the conversation about the nature and direction of science, culture, art, medicine, and technology in our world – going on both within and without the small corner of that world inhabited by philosophers.

Chapter 2

Incentives Arguments for Intellectual Property

The central idea of the incentives argument is simple. Because it is so easy to copy images, songs, novels, designs, and formulas, and because it is so hard to prevent individuals from copying these once they are out in the world, there is little incentive to invest in the work necessary to create these things. But granting IPRs to inventors, researchers, and creators of literary and artistic works protects them from copying, and provides a way for them to make a living by selling the products of their creative and innovative work. Thus, to incentivize this work, and ensure a sufficient supply of good ideas, we should grant IPRs.

Incentives arguments for IP have come under relentless attack from those who hold that IPRs are inimical to the task for which they were intended. In the age of rapidly evolving technology for copying, altering, sequencing, tweaking, and generally accessing everything from media to genes to hard-to-get mid-20th century jazz recordings, IPRs are no longer a spur, but rather a barrier, to innovation. Patents and copyrights are far stronger than is necessary to incentivize intellectual labor; they overcompensate their holders, and allow big, multinational media companies and pharmaceutical firms to quash creative work by individual artists and garage tinkerers. IP is an anachronism; patent and copyright are 19th century institutions that regulate 21st century industries, and it is time for them to go – or at least, undergo serious renovation.

Sentiments like these are widely voiced in IP theory, and widely considered a serious challenge to incentives arguments for IP.²⁶ Sometimes arguments for these points take the form of careful, sober analysis of the complicated economic, legal, and policy

²⁶ For representative examples, see Angell 2005; Benkler 1999; Bessen and Meurer 2009; Boldrin and Levine 2008; Boyle 2003, 2010; Heller and Eisenberg 1998; Kinsella 2008; Lessig 2002, 2005; Shiffrin 2008; Wilson 2006.

issues involved, and sometimes they take the form of prophetic warnings about the folly of patent lawyers and the dangers of standing a’thwart the progress of science, technology, and culture.²⁷ Running through both sober analyses and blazing jeremiads are two claims that are contra the main thrust of incentives arguments: IPRs are not really necessary to incentivize creative and innovative work, and IPRs are so strong they overcompensate their holders and do more harm than good.

This chapter is neither a defense of incentives arguments for IP, nor is it a criticism of anti-incentives arguments. In large part both of these rest on claims about economics, law, and politics that can only be settled by determining the actual economic, legal, and political facts. Rather, this chapter is about the dialectical relationship between the pro and con sides of the incentives argument. In this chapter I argue that the popular criticisms of the incentives argument fail to appreciate three important elements of the argument, and thus fail to offer a convincing rebuttal to it. IPRs don’t just incentivize intellectual labor; they also incentivize production of certain kinds of goods, incentivize investment in certain kinds of industries, and finance science, technology, and culture through a certain kind of mechanism – namely, a market. Criticisms of IP often fail to address these other dimensions of the incentives structure created by IP institutions.

The conclusion I draw from this discussion is that we must recast arguments about incentives as involving a philosophical problem about what kind of creative and innovative work we want, what kinds of goods we want markets for, and how we want to pay for them. It is not enough to ask whether or not IPRs are efficient mechanisms for providing the incentives we desire; we also have to weigh the relative advantages and

²⁷ See, for instance, Barlow 1994.

disadvantages of incentivizing certain kinds of work and production over others, and the relative advantages of IPRs as a mechanism for creating incentives structures for regulating research, development, and creative and artistic work over other mechanisms, such as direct public financing. I begin in the next section with a brief overview of the unadorned incentives argument, and then lay out two recent, sophisticated presentations of the main strands of the anti-incentives arguments in section 2. I criticize these two presentations (and through them, the main ideas behind the anti-incentives argument) in sections 3 and 4, and draw some conclusions in section 5.

1. The Logic of the Incentives Argument

As we discussed in the introduction, ideas are commonly held to be both nonrival and nonexcludable. For the paradigmatic subjects of old-fashioned property rights – namely, land and material objects – there are costs to use, and drawing physical or legal boundaries around them to exclude users is not (at least, not always) terribly difficult. This is not the case with the subjects of IPRs. If I build an engine based on your design, produce a drug, download a digital copy of a song or novel, or make a batch of Coca-Cola in my bathtub, anyone else is still in a position to do all these things with the same ideas I used. My use of these ideas does not deplete the amount of them that's left in the world.

Another way to put it is, once the cost of making the good available to one person (the production cost of the good) has been paid, the cost of adding additional users does not add to the total cost (Landes and Posner 2003, 20). The cost of one hundred users, in other words, is the same as the cost of one. Thus, for any idea (say, a novel), there are,

for all practical purposes, indefinitely many possible tokens of it. Even if every extant copy of a novel is being used, there's nothing that prevents anyone from making a new copy so that you can read it too.

This is, remember, for the *idea* itself. The marginal cost of production for *tokens* of the idea is not zero, and will fluctuate. The marginal cost of production of tokens of *Moby-Dick*, for instance, is higher for printed books than for digital copies of the novel. The cost of a digital copy of the book is practically zero (there is some opportunity cost in time, and there is a small cost in use of bandwidth).

Take any of the above examples – a design for a drug or engine, copy of a novel, or formula for Coca-Cola. In some form, if I don't want anyone to use these ideas, it's possible for me to exclude other from them. So, if I write the secret formula for Coca-Cola down and lock it in a safe, and only take it out to make myself a glass every now and again, its possible to keep the ideas exclusive. But suppose I make some Coca-Cola for a friend, who is adept at chemistry. My friend figures out the composition of the drink and reverse engineers the formula, then starts making it for herself. It's very difficult for me to stop my friend from doing so; once I've made it available to her, there's not much I can do (short of force) to prevent her from making the drink at will. When realized as a recipe on a piece of paper locked in a safe, my formula is fairly excludable. But when realized as the drink it is a recipe for, it becomes difficult for me to exclude others from using it. The same goes for certain realizations of other ideas: digital files of songs, or schematics in a patent application, or any realization of the ideas that are easily reverse engineered.

Subjects of IPRs, such as songs, are thus (for all practical purposes) nonexcludable: once they have been made available to one person, it's difficult to prevent them from being available, in principle, to *anyone*. Songs, for instance, are easily copied, and, in the form of digital files, easily accessed. Once a song has been made available to one person (say, in the form of a CD), it is easy to turn it into a digital file, add it to one's library, and share digital copies of it on a peer-to-peer file-sharing site.

In virtue of having these two features, ideas are commonly considered *public goods*. Because ideas are public goods, markets in ideas (allegedly) will not be efficient. Because ideas are both nonrival and nonexcludable, the cost of using them is very small, and it's very difficult to exclude anyone from using them, once the production cost of the ideas has been paid and they have been made available to anyone. The production cost of ideas, however, is not negligible, but often significant; consider the price of building, supplying, and staffing a lab capable of doing serious biomedical research, or the costs in time, energy, and resources of writing a novel or producing a record. Without a means to exclude others from use, it is difficult to benefit from the production of ideas, and difficult to recoup any production costs. But, once the idea has been made available to any person, anyone can access it for very little cost. Thus there is little incentive to produce ideas, because production is so costly, but there is great incentive to wait around for others to produce, because *use* is not costly. Given this incentive structure, a market in ideas will (allegedly) fail to produce an adequate supply to meet demand, because production of ideas is so costly and has little (financial) benefit.

In short, because as many copies or tokens of ideas as the market can bear can be produced without depleting or subtracting from the value of ideas, and because it is very

difficult to prevent copying of ideas once they are out in the open, the price of a novel, drug, engine, and the like will very quickly approach the marginal cost of production. This is – so the incentives argument goes – likely to be much lower than the cost of producing the good, and so there will be little incentive to invest in creative and innovative work, and the production of goods like drugs or novels. Supplying ideas, in other words, presents us with a classic *public goods problem*. Everyone benefits from having a certain supply of good ideas and things like novels, drugs, and engines. But because ideas are both nonrival and nonexcludable, it is very difficult to get everyone who benefits to internalize a share of the cost of provision. The result – so the argument goes – is that a market in goods that are tokens of ideas (printed copies of novels, say, or CDs) will not provide sufficient return on investment in labor and capital to incentivize intellectual labor. Markets for novels, drugs, CDs, and the like – sans IPRs – thus fail to ensure an adequate supply of good ideas.

Granting IPRs, however, can correct market failure in this case. IPRs allow individuals to exclude others from using the ideas they've produced, and thus provide a mechanism to recoup production costs and benefit financially from producing (by charging for use and access), thus creating an incentive to produce ideas. In reality, of course, IPRs can do little to prevent, ahead of time, unauthorized use. The prevalence of copyright piracy today shows that they are not terribly effective as preventative measures. IPRs can protect their holders against unauthorized use, though, in two ways. First, they can act as a deterrent, by setting penalties for infringement. There's good reason to believe, for instance, that the possibility of costly patent infringement litigation is a fairly

effective deterrent.²⁸ Second, IPRs provide a means for patent and copyright holders to seek redress from unauthorized users, through the intermediary of the courts.

All of this is deeply familiar to anyone even somewhat versed in IP theory. The basic idea behind the incentives argument is simple: without IPRs, there will be a market failure for ideas, which will result in an undersupply of (good) ideas relative to demand. IPRs, by allowing holders to exclude others from use of ideas, create incentives to produce. The *purpose* of IP, according to an unadorned version of the basic incentives argument, is to correct this market failure: to turn ideas from public into excludable goods.

Let me restate the same idea in a different, slightly more rigorous, form. Consider a parcel of land held in common by a group of individuals, where each individual has privileges of use and access in the land. Each of these individuals can use the land as much or as little as they want, but no individual has the ability to exclude others from use, or any exclusive claim to any part of it. There is a maximum level of use, L , that can be made of this resource. The optimal individual share of this resource is L/n , where n is the number of individuals in the group. However, if any individual uses *more* of the resource than the rest, then the share of the resource available to the other individuals will decrease. In this situation, there is no way to prevent any individual from using more than L/n , and really no reason for any individual to stick to that share. The dominant strategy for any individual user is to consume as much as she can as quickly as possible; otherwise, others' use may preclude her from getting anything at all. When each individual has a common

²⁸ Perhaps *too* effective. There are worries that patents deter not only would-be infringers, but raise the costs of accidental infringement so high, they also disincentivize legitimate research (Heller and Eisenberg 1998).

privilege, but no mechanism exists to prevent any individual from overuse, L will be reached very quickly, and the resource will be depleted. This is the familiar tragedy of the commons.²⁹

This situation is, for well-rehearsed reasons, unsustainable.³⁰ In order to ensure that the situation does not turn tragic, there is a need for some public and enforceable set of rules about use and access to the good. These rules could be of different kinds; for instance, a steep tax on any individual who consumes more than their share, which will disincentivize overuse and compensate others for any one individual's shenanigans. Property institutions, which give certain individuals rights of use and access to the good, are another way to avert the tragedy of the commons.³¹ Property institutions here have some advantages over other ways of averting a tragedy. First, by allowing individuals to exclude others from use of their shares, property rights ensure that individuals internalize any benefits of productive use of these shares. Since individuals are excluded from use of any other portions of the land but their own, they also have to internalize the costs of overuse of their plot of land, thus incentivizing responsible management of the land. In sum, property rights, by virtue of forcing individuals to internalize the costs of

²⁹ This models the tragedy of the commons as a prisoner's dilemma; for discussion see Ostrom 1990, 3-5.

³⁰ "...therefore if any two men desire the same thing, which nevertheless they cannot both enjoy, they become enemies; and in the way to their End...endeavor to destroy, or subdue one an other. And from hence it comes to passe, that where an Invader hath no more to feare, than an other mans single power; if one plant, sow, build, or possesse a convenient Seat, others may probably be expected to come prepared with forces united, to dispossesse, and deprive him, not only of the fruit of his labour, but also of his life, or liberty." Hobbes, *Leviathan* Ch. 13 (Hobbes 1996).

³¹ Grouping both liability and property rules together as strategies for governing use and access comes from Calabresi and Melamed 1972.

overconsumption and allowing individuals to internalize the benefits of productive use, disincentive the former and incentivize the latter.

Now suppose the good in question was not a parcel of land, but a design for a new internal combustion engine. Suppose a group of individuals with sufficient technical expertise and heterogeneous preferences all have common privileges of use in the design. Some will want to build the engine, some will want to tinker with the design to see if they can improve it, some will want to set up a business building engines to sell to their friends and neighbors, and some may just want to see if they can figure out how the engine works, and admire its inventor's ingenuity.

Unlike with land, use and access does not diminish the value of the design, and so does not add to the cost of using the design. No matter how many times the design is tinkered with, or used to produce an engine, the opportunities for the same (and many other) uses do not diminish, and this holds even if additional users of the design are added. Each *use* of the design, however, has a cost. Some of these uses (admiring the design's ingenuity) are not all that costly; some (building engines) are very costly. Each user has only a finite amount of resources to invest in using the design and in advancing any other goals they may have, and must decide what to do with these resources based on what investments of it are most valuable to them.

Suppose there is a subset of the n users who want to use the design to build engines, and have the expertise to do so. Some of these want to build engines for personal use, some want to build engines to sell to others for profit. Using the design to build engines has a certain cost. Because the design is nonexcludable, no individual can exclude others from making engines, and because it is nonrival, building engines does not

preclude any other user from doing so in the future. Any user who starts selling engines can reasonably assume that they will have to continue dropping prices to compete with further users who make and sell the engine, as the market becomes flooded with the good. Thus no user can be assured that they will make enough off selling engines to cover their initial investment, or generate enough revenue to make the investment worth it (compared to other possible uses of their resources). Those who make the engines for sale will not do so, and those who want the engine (but lack the technical expertise to make one for themselves) will not get it.

Granting IPRs, however, allows the holder of the patent on the design to exclude others from using it to make engines. They are now assured that they can either charge enough, or license their patent for enough, to make their initial investment worth it. By allowing holders to exclude other users, IPRs thus allow their holders to internalize the benefits of certain uses of ideas, which incentivizes these uses. This should ensure that there is a sufficient supply of certain goods (engines, drugs, mystery novels, smartphones) to meet demand for them.

The second form of the argument makes the same basic point as the first, but the second form allows us to see a few important things. First, both IPRs and old-fashioned property rights function similarly in creating an incentive structure by distributing costs and benefits of activity. From the perspective of one theory of property – that property is an institution that allows individuals to internalize externalities – there is no conceptual difference between IPRs and old-fashioned property rights.³² Second, overuse is not a problem for nonrival goods like ideas. The only reason to exclude some from use is to

³² This conception of property comes from Demsetz 1967. The application of Demsetz's view to IP is discussed in Lemley 2005 and Duffy 2005.

incentivize certain kinds of uses. The more ideas are used, the more potential value they can generate. Since we don't have to worry about overuse and depletion, we should want individuals to use ideas as much as possible. Third, though overuse is not a problem, *underutilization* is. This is the whole point of IPRs in the first place, to create incentives so that ideas won't be underused. But IPRs also forestall potentially valuable and productive uses of ideas, by restricting access to and use of ideas. Thus this is also a reason not to make IPRs too strong.

The reasoning behind the incentives argument thus sets two conditions on the justification of IPRs. First, the ideal situation is to have ideas remain in the commons, as much and for as long as possible. IPRs are a necessary evil; they destroy common privileges of use and access in ideas, thus forestalling potentially valuable uses of ideas, in the interest of incentivizing uses that will have greater value. But if they are not necessary for this purpose, then they are just an evil, and should be done away with. Second, IPRs should only be as strong as they need to be to create the incentive structure we desire. This means that they should only allow their holders to internalize as much of the value generated by productive uses of ideas as is necessary to incentivize these uses, and no more. If IPRs are too strong, then they should be altered.

2. Two Criticisms of the Incentives Argument

IP is under attack on both these fronts. There are two general arguments, often rehearsed in anti-IP literature, which hold that IP institutions fail to meet either condition. Here I will examine two prominent and sophisticated versions of each of these general arguments. Though my concern here will be specifically with these two versions of the

arguments, the criticisms they make are ubiquitous within the literature on IP. It's safe to say that these arguments make claims about how incentives arguments for IP fail that are widely shared by many IP theorists.

The first argument comes from a recent book by Kal Raustiala and Christopher Sprigman (Raustiala and Sprigman 2012). Raustiala and Sprigman (for ease, I will refer to them as "R&S") hold that the incentives argument makes an empirical claim; namely, that without giving individuals exclusive rights to certain uses of ideas, these individuals will lack incentives to exercise those uses, and will therefore choose not to do so. This claim, according to R&S, is simply false. There are multiple industries in which there is no IP protection, and what we find in these industries is that individuals still engage in creative and productive activity, and learn to cope with ubiquitous copying without going out of business.³³

R&S attack what they call the "monopoly theory of innovation": "Innovation requires rules that allow creators to control who can make copies – either by making the copies themselves, or selling licenses to others" (Raustiala and Sprigman 2012, 6). The monopoly theory of innovation is an empirical claim: without this control, the incentives necessary to stimulate creative work are absent, and so this work won't happen. R&S's argument revolves around the analysis of three industries: fashion, cuisine, and stand-up comedy. Each of these industries is notable for two reasons. First, all three are businesses in which it seems likely that copying can be severely detrimental to any individual's ability to make money. If I come up with a new design for a dress, for

³³ There are other topics in R&S's book that I am not going to discuss here, such as the use of cultural or community norms for enforcing prohibitions against copying, and their relative advantages over legally enforceable IPRs.

instance, and Wal-Mart copies the design and sells dresses for cheap in all of their stores, it will have a negative effect on my business. Second, in all three, copyright and patent protections were, due to various historical circumstances, never extended to cover the ideas upon which their products are based – namely, clothing designs, recipes, and jokes.³⁴

If the monopoly theory of innovation is correct, we should expect to find severely depressed creative work in these areas, at least relative to a baseline of creativity we could expect with IP protections in place. But of course, we find the exact opposite. All three are currently thriving industries; the apparel industry, for instance, generates an estimated US\$1.3 trillion worth of goods annually (Raustiala and Sprigman 2012, 22). This is not to say that copying is an acceptable practice within these industries. Rather, what they claim is that fashion designers, chefs, and comedians have found ways to cope with copying, without legal protection against it. Within the community of stand-up comedians, for example, there are tacit norms governing when and under what circumstances comedians can borrow other comedian's material, and which assign enforceable penalties for violations (Raustiala and Sprigman 2012, 105-115).

R&S thus conclude that the monopoly theory of innovation fails. This not to say that IPRs are *never* necessary or desirable; they accept that there are some industries, such as the pharmaceutical industry, which seem unable to get by without them. What it means is that it is not the case that, as a general rule, exclusive rights to ideas are necessary to incentivize use of ideas. In some industries, the gains from unrestricted release of nonexcludable goods into the market are sufficient to provide incentives, and

³⁴ For these respective histories, see Raustiala and Sprigman 2012, 22-36, 59-69, 103-115.

IPRs are superfluous (for providing incentives). Creative individuals will, in short, find creative ways to make a living off creative and innovative work. In many ways, what IPRs actually do is stifle creativity, by making it unnecessary to innovate ways to make and market valuable uses of ideas. By forcing individuals to figure out how to make their business work in the absence of legal protection against copying, the lack of IPRs in these industries may actually stimulate creative work. This last bit is a widely echoed sentiment in IP theory; one sees it, for instance, in concerns about how IP protections will stifle creative work in biotechnology.³⁵

The second argument comes from Mark Lemley (Lemley 2005). Lemley argues that property rights in ideas allow their holders to internalize far more of the value generated by productive uses of ideas than is necessary to incentivize those uses. Lemley makes a distinction between what he calls the “average fixed costs” of using ideas, and the “total social value” of an idea (Lemley 2005, 1053-1055). The first is the cost of generating a commercially viable product from a design, formula, manuscript, or the like.³⁶ The second is the total value generated by the ideas; not just the returns on sales of commercially viable products, but also the use-value to consumers of these products, and the positive externalities generated by uses of the ideas and products developed from them.

In order to incentivize uses of ideas, IPRs must allow users to get sufficient returns to cover the average fixed costs of production plus a margin that makes initial investment in production worthwhile. Anything beyond that is too strong, and thus

³⁵ See, for instance, Carlson 2011.

³⁶ Lemley refers to these as “informational goods” (Lemley 2005, 1053).

overcompensates producers. There are significant drawbacks to overly strong IPRs, and I want to highlight one Lemley mentions in particular.³⁷ Overly strong IPRs create dynamic inefficiencies (Lemley 2005, 1060-1062). Innovation requires building on preexisting ideas, to create new ones. IPRs give their holders the ability to control further uses of protected ideas. This has a chilling effect on downstream uses of these ideas; either IPR holders can prevent these downstream uses, or they can require compensation for them, thus raising the transaction costs of downstream innovation. The result is a dynamic inefficiency: short-term gain (in the form of increased productivity in certain industries) is prioritized over long-term welfare (in terms of efficient utilization of ideas).

The upshot of this is that overly strong IPRs allow their holders to pocket far more of the total social value of ideas than is necessary to provide incentives, and at significant costs to others. By giving their holders gatekeeping power over further uses of protected ideas, IPRs give their holders the ability to pocket some of the gains from downstream uses of their ideas, and thus internalize some of the positive externalities of downstream innovation. This not only reduces the value to potential downstream innovators (by raising the transaction costs of downstream work) but also raises the costs to consumers (by raising production costs, and thus forcing prices up for products developed out of downstream innovation).

Lemley is somewhat agnostic in his article about whether actual existing IP institutions overcompensate IPR holders. He quotes with approval the remark by Fritz Machlup from the preface to this essay, that he (Machlup) could not in good conscience

³⁷ The other four are: IPRs create static inefficiencies in the form of deadweight losses (I discuss this in section 6), encourage socially wasteful rent seeking, their enforcement imposes administrative costs, and can lead to distortionary overinvestment.

recommend creating a patent system if it did not exist, or eliminating the one that we already have (Lemley 2005, 1065). Lemley's concern is that by treating IPRs as on a par with old-fashioned property rights, overcompensation of the kind just detailed will occur. Doing so is costly; no one can possibly be entitled to, let alone require as incentive, so much of the total social value of ideas. It is akin, according to Lemley, to holding that individuals who admire a garden in front of a house on their evening walk have to compensate the homeowner for the labor she put into planting the flowers. His claim is that we should treat IPRs not as property rights, but as a form of government subsidy, meant to aid certain industries. The upshot of this is that, for Lemley, the incentives argument may give reasons for certain kinds of instruments, which may or may not look like existing patents and copyrights, but does not give a justification for intellectual *property* rights.

3. Expanding the Incentives Argument from the “Demand” Side

Both of these arguments fail. That they fail is significant, as the basic ideas of each are widespread within IP theory. But it is less interesting *that* they fail than it is *how* they fail. Each argument expresses some legitimate concerns about IP institutions. However, in presenting these concerns, each argument makes some important conceptual errors. These errors tell us a great deal about how ideas *qua* subjects of property, and IP institutions, are thought of within IP theory and policy. But they also tell us a great deal about how the incentives argument should be interpreted. I will take each argument in turn.

To begin, the argument from R&S. R&S fail to show that the incentives argument does not justify IP. This is because none of the three industries they discuss is a sufficient counterexample to the incentives argument. The incentives argument holds that, insofar as ideas are *nonrival* and *nonexcludable*, there is a lack of incentives for certain productive uses of them. Even if ideas are nonexcludable, but some *tokens* of these ideas *may be* excludable. Anything that adds value to the token thus makes a valuable product, even if the ideas realized by the token can be copied freely. Indefinitely many tokens of *Moby-Dick* can be made, but there are only so many first editions of the book. These editions are valuable, no matter how many copies of the book there are or could be.

Each of the three industries is based around using ideas to produce tokens which are valuable not *just* because they are realizations of those ideas, but because of other factors that cannot be so easily copied. People do not pay for clothes just because they like the way they look. They also pay for certain brands, they pay extra to shop in certain stores, they buy clothes because celebrities or friends are wearing similar items, and they buy clothes that project a certain image. People will pay extra for clothes because they are in a certain store, have a certain label, and are associated with certain people; these factors cannot be copied along with clothing designs. The same with comedy and food. People do not pay to see stand-up comedians just because the jokes they tell are funny, they pay to see them because the comedian's *performance* of the jokes is funny. People pay to go to restaurants not just because of the food, but because of factors such as the ambience, the social aspects of dining out, and the sheer and exquisite pleasure that comes from having someone else cook and bring you food and wine.

R&S admit this, when they discuss the distinction between “product” and “performance”:

...the central item created by a chef – food – cannot be easily disentangled from the “packaging” of the restaurant in the way that songs can be separated out and sold, or traded, as discrete digital files...There are ambience, service, energy, and other intangibles in the mix. All of these factors work together. Copying one aspect – the main dish – may be easy. Copying the experience in full is virtually impossible. The experience is less of one buying a product and more that of enjoying a performance. (Raustiala and Sprigman 2012, 85-86)

R&S believe that this is a big reason why creativity flourishes in the culinary and comedy industries, even without IP protection. They fail to see the negative consequences of this for their argument. What their argument shows is not that the incentives argument fails because fashion, food, and comedy show that IPRs are not necessary to provide incentives; rather, their argument shows that the incentives argument *doesn't apply* to these industries because the incentives for creative work in these three industries are already in place. What makes a song valuable is the song; the packaging matters fairly little, and then only to people with certain kinds of preferences, which is why the nonexcludability of recordings of songs is a problem for the music industry. But what makes comedy, food, and fashion valuable is more the “packaging” than anything else. If the author of this essay humbly offered to cook you (the reader) a molecule-for-molecule duplicate of your favorite dish from your favorite restaurant, it would in no way substitute for dinner at that restaurant. Jokes, designs, and recipes may be nonexcludable, but designer clothing, tables at nice restaurants, and tickets to comedy clubs are *highly* excludable, and as long as these industries revolve around generating revenue by selling *these* excludable goods, the lack of legal protection against copying is less of a drag on business.

In short, the incentives argument applies when tokens of ideas are valuable *because they are tokens of ideas*. CD's, paperbacks, photographs, and newspapers are valuable mostly because they are media that deliver songs, novels, images, and information. These industries are struggling because they cannot package these goods in forms that are valuable in and of themselves, as opposed to just delivery systems for the ideas. But handbag designs, jokes, and recipes are valuable because they produce goods that come in an attractive package; these are valuable not just because they are tokens of ideas. In these cases the incentives argument doesn't apply; the incentives are already there. But in cases like the former, the incentives are missing; in these cases, IPRs may be a necessary and efficient tool for providing the necessary incentives.

This is an important point about the incentives argument, one that is often overlooked in IP theory. Discussions of the incentives argument are almost always from the "supply" side. From the "supply" side, the point of IPRs is to incentivize creative and intellectual work, and uses of ideas. However, there is also a "demand" side to the incentives argument. From the "demand" side, the point of IPRs is to incentivize *production of desirable goods*. In other words, the point of IP is not just to incentivize the production of manuscripts, formulas, and blueprints, but also the production of novels, medicines, and engines. Sometimes IPRs are not necessary to incentivize the production of desirable and valuable goods; copyright protection of recipes, for instance, is not necessary to incentivize building and opening restaurants.

We need, therefore, to modify the basic version of the incentives argument. Because ideas are nonexcludable, there is little incentive to invest labor and capital in their production when there are no other factors influencing the value of tokens of these

ideas, other than that they are tokens of ideas. Creative work that can yield products that are valuable as products (and not just as tokens of ideas) will be able to attract investment no matter what; this is why people can still get loans to open restaurants, comedy clubs, and boutiques without copyright protection for jokes, recipes, and designs. This has potential costs. Consider a set of ideas in which a group of n individuals has common privileges of use and access. They can use these ideas to make two kinds of goods, call them a goods and b goods. Each producer could make both a goods and b goods, and each has finite resources with which to do so. Let us say that, for many consumers, a goods and b goods are substitutable. However, for a significant group of consumers, they are not substitutable. Because this is the case, demand for b goods will depend on the availability of units of a goods; many consumers will just buy a goods no matter what (since given their preferences the two goods are substitutable), so demand for b goods will be negatively affected by the level of availability of a goods. A sufficient level of availability of a goods will thus reduce demand for b goods to the point at which returns for producing b goods will no longer make their production worth it.

Now suppose that demand for a goods will stay relatively constant, while demand for b goods will diminish quickly after they are introduced to market because b goods are easily copied. Each producer will invest in producing a goods, and produce only a few b goods, and only produce b goods up to the point at which market price stays above marginal cost of a b good. The result is that consumers who prefer a goods, or who are neutral between a goods and b goods, will have sufficient options to satisfy their preferences, but not those who prefer b goods, because producers would either have made only a few b goods or none at all.

In short, widespread copying effectively shifts investment towards the production of certain kinds of goods. This is a clearly observable effect; it is why getting an mp3 version of a song or album is so easy and cheap, but buying a CD is such a hassle. It is also why getting a standard pair of brown men's loafers costs you around US\$20 and a trip to Payless, but getting a slightly more unique pair of shoes can cost hundreds of dollars. This is, in and of itself, not a bad thing; there are no claims here about what we *should* do. The point is twofold: first, the incentives argument is not just about incentivizing production and use of ideas, it's also about incentivizing production of certain *kinds* of goods that realize those ideas. Second, rebuttals to the incentives argument have to show not just that IPRs are not needed to incentivize creative and intellectual work; they have to show that they are not needed to incentivize development of the products that we want, or at least that we need. They thus have to address ranking of different sorts of goods and preferences, and explain why some goods are either better than others, or some preferences or needs are more important than others. This is tantamount to an appraisal of how our art, science, culture, technology, and medicine *should* be, which goes beyond purely technical questions about incentives.

4. Investment and Downstream Creative and Innovative Work

Lack of attention to the “demand” side of the incentives argument is also a problem with Lemley's argument that IPRs risk creating dynamic inefficiencies. One way in which a dynamic inefficiency can be created is if downstream innovation is chilled by IP protections. But another way is if lack of IP protection stimulates overinvestment in

certain kinds of creative and innovative work, and the development of certain kinds of products.

Suppose producers can choose to use an idea to produce units of a very valuable, but easily copied product (c goods), or tinker to get new ideas which can serve as the basis for other, different products (d goods). Suppose that both are easily copied, and that there is no protection against copying. Tinkering to make new products takes a long time, and has many risks, as there's no guarantee that d goods will be valuable at all. However, c goods are very valuable, and production and sale of c goods will generate immediate, guaranteed returns. Further, since there is only a finite demand for c goods, the more the ideas are used to make c goods, the quicker use of these ideas becomes unprofitable. The result is that, as some producers introduce c goods, others will shift resources away from tinkering to the production of c goods, to grab as much market share as possible, and c goods will flood the market very quickly. As long as c goods are valuable, there is no incentive to risk capital in the production of d goods. The ideas have been underutilized; a potentially value-creating use of these ideas (production of d goods) never happened, in the service of generating guaranteed short-term returns by making c goods, resulting in a dynamic inefficiency.

This illustrates another important point about the incentives argument, one that, again, is unfortunately often overlooked in IP theory. The point of IPRs is not just to incentivize production of ideas and products that realize ideas; it's also to attract *investment* to certain kinds of creative and innovative work, and development of certain kinds of products. Even if it's true that people will still do creative, scientific, and inventive work in the absence of IPRs, and that business will find a way to cope with lack

of IP protection, the lack of IPRs can still have a chilling effect on investment in certain kinds of creative work, and certain kinds of R&D. Intellectual labor and product development, after all, require capital, and capital will flow towards the point at which it can yield the highest returns.

The inefficiency in the above scenario was caused in large part because one use of ideas (for short term gain by producing *c* goods) is a better investment than another. As long as any producer can copy *c* goods, there is no incentive to invest in downstream innovation to get *d* goods. But if *c* and *d* goods are protected from copying, then there is incentive to invest in *d* goods, first because there is only a finite amount of capital that can be sunk in *c* goods before returns to investment diminish, and second because the risk in producing *d* goods is now less of a drag on attracting capital, because of diminished returns of investing in *c* goods, and because of increased value of potential *d* goods due to restrictions on copying. Contrary to Lemley's argument, IPRs (in this scenario) help incentivize downstream innovation because they incentivize *investment* in it.

There is no part in the literature on IP where this needs to be appreciated more than in discussions about reform of pharmaceutical patents.³⁸ Critics of the pharmaceutical industry regularly invoke the wide profit margins pharmaceutical companies can command on sales of their most profitable drugs, as evidence that pharmaceutical patents are far stronger than what's necessary to incentivize R&D. What is not often discussed in these criticisms is that the pharmaceutical industry is made up not just of firms that develop drugs, but also firms that exclusively make generic versions of off-patent medicines. Investment in the pharmaceutical industry can flow into either

³⁸ For examples of criticisms that overlook the following point, see Angell 2005; Flory and Kitcher 2004; Pogge 2007, 2008; Reiss and Kitcher 2009.

kind of company. What would happen to investment in R&D if there were lighter restrictions on making generics? It's reasonable to think that the pharmaceutical market could come very quickly to resemble the market in *c* goods and *d* goods. That is, investment could shift away from R&D of *d* goods (new medicines) and towards production of *c* goods (versions of existing and proven drugs), because the immediate returns in making and selling existing, proven medicines make them a more attractive investment option, given the risk and cost involved with developing new drugs.³⁹

The main thrust of Lemley's argument, however, concerns the disparity between what is needed to incentivize production and the returns that IPRs can provide their holders, and its against this that the above point about IPRs incentivizing not only production but also investment is the most damning. Perhaps covering average fixed costs and a reasonable margin is enough to incentivize intellectual laborers to do their thing, but is it enough to incentivize investors to sink capital into pharmaceutical companies and small biotech firms? Could US pharmaceutical companies really attract enough investment to spend (roughly) US\$50 billion on in-house R&D if the only returns they could show investors were "average fixed costs plus a reasonable margin"?⁴⁰ There

³⁹ Its possible to view the phenomenon of "me too" drugs – that is, drugs that duplicate the functions of existing, proven medicines, but are chemically different and so not covered by these existing drugs' patents – as a version of this scenario. Drug companies, recognizing that there is demand for drugs that perform a certain function, invest great deals of resources in getting their own versions of drugs that have these functions, as well as marketing efforts to grab market share from competitors. The production of these "me too" drugs thus may show a tendency, within this industry, to send capital down the path of least resistance. Of course, the existence of "me too" drugs also shows that patent protection can only go so far in forestalling this kind of thing.

⁴⁰ The number is from Reiss and Kitcher 2009. This number is contested; how you count depends on what you consider "R&D". But, for what its worth, the Pharmaceutical Research and Manufacturer's Association (the infamous "Big Pharma"), the trade group

are many options for investment; what may seem to us like ridiculous returns may very well be what is necessary to keep capital flowing into industries that rely on IP protection, instead of other businesses.⁴¹

This kind of argument could easily be, for Lemley, a *modus tollens* to my *modus ponens*. To Lemley this may be like arguing that we should pay homeowners to plant the gardens we enjoy, because its more enjoyable to take walks in neighborhoods with nice gardens than those without. Another way to take the above is that IP sets up a system in which the only way to finance intellectual labor is to allow laborers and investors to capture a huge portion of the total social value of the ideas they create and develop. This is an unacceptable situation, and is as much a reason for IP reform as it is a defense of IPRs as incentivizing mechanisms.

This brings us to a third important point about the incentives argument. IPRs do not incentivize intellectual laborers by allowing them to internalize a share of the “total social value” of ideas. Rather, IPRs incentivize intellectual labor by giving laborers *the chance* to internalize a share of the *economic* value of ideas *indirectly*, by facilitating the creation of *markets* in products developed from those ideas. The total social value of ideas is likely far beyond what can be obtained from selling tokens that realize those ideas in a market. The value of an IPR is equivalent to the financial return on selling tokens of protected ideas (or licensing the IPR). The price of these goods tracks desire, and desire tracks perception of value, not value. A world-changing, carbon-neutral

representing the pharmaceutical industry, gives the same number as Reiss and Kitcher. They report in-house R&D spending for 2011 (the last year of reported data) as \$US48.6 billion; see “Industry Profile 2013”, at <http://phrma.org/sites/default/files/pdf/PhRMA%20Profile%202013.pdf>.

⁴¹ John Duffy makes a similar point in his comment on Lemley’s article (Duffy 2005).

electric car may well fetch far less money than we think it should, because people would still rather buy SUVs that get eight mpg on the highway. Further, consumers have finite resources at their disposal, and have to make decisions about what to buy knowing they cannot satisfy all their preferences. In a market where consumers have heterogeneous preferences, revenue will fall short of total value, even if consumers have perfect information. Older individuals who are aware of the dangers posed by climate change may discount benefits to future generations from reduced atmospheric carbon, and buy SUVs anyway. Financial returns on IPRs are likely always going to be only a share of the total value of the ideas, because market price for tokens of ideas will not reflect the value of externalities generated from just having the ideas.

In fact, when compared with Lemley's standard of overcompensation – that is, capturing the total social value of an idea – one could make a case that IPRs, if anything, *undercompensate* intellectual laborers. Imagine that, instead of rewarding the inventor of penicillin or AZT or a cheap, world-changing, carbon-neutral electric car with a patent, and allowing them to recoup production costs and generate whatever profits they can by selling these items, we instead held an auction.⁴² Suppose that there was a certain price, and if the auctioneer could not get a bid at that price, the good would not exist. This price is unknown to everyone, including the auctioneer (let's say its in a sealed envelope, and she'll open it when the auction is over). So, for every new price offered by the auctioneer, you don't know if failing to bid will result in the good disappearing from the face of the Earth, or whether you've met or surpassed the necessary threshold. Thus, at

⁴² This is inspired by Ronald Dworkin's hypothetical auction among shipwreck survivors (Dworkin 2002).

every stage, you must weigh the cost of paying the auctioneers proposed price, with the cost of not having the good.

Now, suppose all bidders were behind a veil of ignorance.⁴³ That is, in bidding, you did not know whether one day you would contract HIV, and need AZT, or whether you would live long enough to experience first hand some of the devastating effects of climate change. How much would you bid for AZT, not knowing if one day you would need to take the drug to save your life?⁴⁴ No matter how high the auctioneer pushed the price, it would always be rational to add one more dollar, because the cost of not having the drug (namely, you'll be dead) is far greater than the cost of meeting the auctioneer's price. If we measure the total social value in terms of what resources we would be willing to commit to have these goods, then market returns on patents and copyrights are far below their total social value.

5. The Incentives Argument is Not Just About Incentives

This is not to say that there aren't serious concerns about IP. For instance, Lemley states five reasons why overcompensation is troublesome. Above, I criticized one; another of these five reasons is relevant, as a possible direct rebuttal to my argument about dynamic inefficiency. Lemley rehearses an old objection to IP: that because returns on IPRs are potentially so much greater than costs, given that IPRs give their holders a temporary monopoly position in a market, too much capital will shift towards IP-protected

⁴³ Ala Rawls's original position argument (Rawls 1971).

⁴⁴ Zidovudine or azidothymidine (AZT) was the first approved anti-retroviral drug for treatment of HIV/AIDS, and remains to this day one of the most successful drugs in history.

industries and away from other investment opportunities, resulting in a deadweight loss.⁴⁵ This is an old objection, introduced by the economist Arnold Plant (Plant 1934).⁴⁶ Above I argued that a major function of IP is to attract investment to certain industries, which without IP protection for their goods may not be an attractive investment option. But if IPRs generate so much value that it pulls capital away from other options and generates a deadweight loss, then this may be a reason against incentivizing investment in intellectual labor with IP.

My remarks are not meant to be a full-scale defense of the incentives argument. My point here is that assessments of the incentives argument have to take into account three dimensions, not often appreciated in discussions about IP and incentives: (1) IPRs don't just incentivize production of ideas, they incentivize production of certain kinds of goods; (2) IPRs don't just create incentives for intellectual laborers, they create incentives for investment in intellectual labor and production of certain kinds of goods; and (3) IPRs don't just allow holders to internalize a share of the value of their ideas, they allow holders to internalize a share of the value of ideas through the mechanism of markets, and thus finance intellectual labor and investment indirectly.

The addition of these three dimensions shows that assessing incentives arguments, and thus IP institutions and IP policy, is not simply a matter of determining what's

⁴⁵ A deadweight loss is a loss in efficiency due to sub-optimal allocation of a resource. In this case, the argument is that the level of investment in patent-protected industries is sub-optimal, because there would be greater gains to efficiency if some of the invested resources were allocated elsewhere. However, because of the great returns on investment (due to the high profits one can reap by having a virtual monopoly on a patented good), resources flow away from other, more efficient investment options to patent-protected industries, thus resulting in a sub-optimal allocation of resources – a deadweight loss.

⁴⁶ For a contemporary discussion of Plant's argument, see Landes and Posner 2003, 22-23.

needed and what's not to incentivize creative and innovative work. There are other issues involved, and it is impossible to disentangle these from discussions about incentives. The incentives argument, in short, is not just about incentives. Assessing the incentives argument requires us not only to consider whether IPRs create incentives, but also to assess the desirability of certain kinds of goods and products over others, of investment in certain kinds of industries over others, and of indirectly financing things like pharmaceutical R&D through markets instead of directly, through mechanisms such as government subsidy or state-funded technology prizes.⁴⁷

Each of these has advantages and disadvantages over the others. The relative strength of each is a function of the interests of the different stakeholders involved and affected. In assessing incentives arguments for IP, we thus have to determine how best to weigh the gains to individuals who prefer, for instance, certain kinds of products, to the losses to individuals who prefer (or need) other kinds. Assessing incentives arguments for IP is really a matter of assessing the relative advantages and disadvantages of different kinds of incentives *structures*, in terms of the effect these will have on markets, investment, and ultimately how the lives of individuals, who consume popular music and depend on pharmaceutical products to stay alive and healthy, go. Thus the incentives problem for IP is not purely technical issue about whether IPRs are needed or not, and what level of protection is appropriate, to incentivize creative and innovative work. It is

⁴⁷ Using rewards and prizes as alternatives to IPRs is widely discussed in the economic literature on IP, and deserves careful consideration. Shavell and van Ypersele 2001 is the touchstone piece; see also Kremer 1998 and Pogge 2008, Ch. 9 for an application of this idea to pharmaceutical R&D. Rosenberg 2012 is another excellent article on alternatives to patents, though in an earlier article (Rosenberg 2004) Alex Rosenberg gives a spirited defense of IP and the incentives argument (he explains this seeming inconsistency in his 2012 article).

also a problem about what kinds of markets, science, technology, and culture we want, how we want to pay for them, and how the interests, preferences, needs, and rights of the different individuals affected constrain the legal institutions we implement. It is, thus, a problem with significant philosophical dimensions.

Chapter 3

Labor as the Basis for Intellectual Property Rights

Of almost equal popularity with the incentives argument, as a possible justification for IP, are arguments having to do with *labor*.⁴⁸ There is a great deal of intuitive appeal to the idea that artists and innovators *should* get something in return for the value they add through their intellectual labor. Property rights secure this benefit, and further ensure that, to paraphrase Locke, those that “desired the benefit of another’s Pains” are forbidden from actions that prevent my enjoyment of those benefits (Locke 1988, 291). IP seems to do the same; since the products of intellectual labor are so easy to access, copy, and use without consent, IPRs seem necessary to secure the benefits of those products for their creators. If I write a good novel or song, isolate a useful organic compound, or create a new drug or a new machine, it looks like I should be entitled to (at least part of) the benefits derived from it, and should be protected from expropriation by governments or free riders. In other words, what holds for land and material objects should hold, *mutatis mutandis*, for novels, movies, songs, drugs, and machines.

But is there a way to make sense of the intuitive notion that labor is a sound basis for IPRs? Can the kinds of considerations about labor that are used to argue for old-fashioned property rights in land and material objects apply to IP as well? In this chapter, I will argue that the answer to both questions is “yes”. I will present and defend a new view, distinct from existing labor theories of IP, which I will call the *productive*

⁴⁸ These kinds of “Lockean” considerations are often cited in the legal literature on IP (for examples, see Gordon 1993; Yen 1990; and Merges 2011, especially chapters 2 and 6), and are widely discussed in the philosophical literature as well (see Fisher 2001; Hughes 1988; Palmer 1990; Shiffrin 2001, 2007; Moore 1997; Wilson 2009; and Mossoff 2012). Considerations about labor have even been cited by the US Supreme Court, as justification for IPRs. For example, in *Ruckelshaus v. Monsanto*, 467 US 986 (1984), 1002-1003: “This general perception of trade secrets as property is consonant with a notion of ‘property’ that extends beyond land and tangible goods and includes the products of an individual’s ‘labour and invention’ ”.

capacities view, and which can make sense of labor as a basis for IPRs. In the next section, I will cover some preliminaries, talk a little about what it means for something like labor to be the “basis” for IPRs, and then examine some other labor theories of IP. I will present my own, new view in sections 2-3, and argue in section 4 that it provides a strong justification for IPRs.

A caveat: though the arguments for IPRs I consider here are broadly “Lockean”, in that they depend on considerations about labor and the connection between labor and property that are in the Lockean tradition, I do not discuss or offer a full-blown Lockean theory of IP, for two reasons. First, such a view would have to bring in the full apparatus of Lockean theory, including discussion of the famous provisos against over-appropriation and spoilage, and would have to discuss the well-known objections to a distinctly Lockean theory of IP.⁴⁹ Second, as I will detail in section 2, there are important differences between the kind of labor theory I offer here and a Lockean theory of IP. The goal here is to give a justificatory argument for IPRs on the basis of labor, not necessarily to give a full-blown Lockean theory of IP.

1. What Does a Labor Theory Have to Show?

By a “labor theory” I mean any view in which labor is either the sole or the dominant *moral basis* for property. The moral bases of property are, according to Lawrence Becker, “...a set of facts about the human condition: facts about human needs, propensities, and behavior from which (together with judgments about values, duties, and virtues) moral

⁴⁹ For these objections and difficulties, see Shiffrin 2001; Hettinger 1989; Wilson 2009; and Hughes 1988.

arguments for and against property rights can be built up” (Becker 1980, 188). Becker makes a distinction between, on the one hand, the moral basis of property, and moral *arguments* for property. There are multiple arguments for property that can be made on the basis of labor: labor-mixing arguments, arguments from self-ownership, and desert-for-labor arguments, to name a few.⁵⁰

Any labor theory has to solve two problems, in order to justify property. First, any labor theory has to explain exactly what it is that is special about labor, which makes it the basis for a claim to property. Call this the *distinctiveness* problem. One way to state the distinctiveness problem is Robert Nozick’s famous remark about pouring tomato juice into the sea (Nozick 1977). Lockean-style labor-mixing arguments hold that labor justifies property because, in laboring, I mix something I have a right over (my labor) with something I don’t (items I labor on). In order to respect my rights over my labor, we have to recognize rights over those items. Nozick’s remark, that laboring may well just be a way to lose these rights (over my labor), rather than gain rights over what I labor on (in the same way that pouring tomato juice into the ocean is a way to lose rights over the tomato juice, not gain rights to the ocean), is a challenge to the idea that mixing what I have a right to with what I don’t is what is distinctive about labor.

⁵⁰ There are a lot of questions that can (and should) be asked about this way of individuating labor theories of property. For example, we need to give an account of what it would mean for labor to be the dominant moral basis for property, in a view that cites multiple moral bases, such as a desert-for-labor view. Becker has devoted a great deal of work to these issues (Becker 1980, 1981, 1992), but these topics are beyond the scope of this chapter. It is sufficient for our purposes here to understand dominance loosely; a labor theory is one in which the preeminent considerations in favor of property have to do with labor, and in which any other moral bases (if there are any) are offered either to support or complement considerations about labor.

Second, a labor theory has to explain why it is *property* rights creators and inventors are entitled to. To see this, consider a desert-for-labor argument for IPRs.⁵¹ On such a view, laborers *deserve* benefit from the creation of some work, because it was their labor that brought it about. Labor justifies property rights because granting property rights ensures that creators and inventors can secure a set of benefits that they deserve. Supposing that laborers do deserve some benefit on the basis of their labor, why is it the *particular* benefits secured by *property rights* that they deserve? Why not, as Becker (1993) argues, offer prizes, or build statues, or give tax rebates, to those who bring about something of value? Why is it the benefits that come from property rights, and not prize money, statues, or passes to the front of the line at the DMV, that are called for?

In short, the idea that labor establishes a desert claim is not sufficient justification for property rights, because (a) its not clear why it is the particular benefits that are generated by property rights, and not some other set of benefits, that are deserved, and (b) for any benefit named there likely is a plausible alternative mechanism for providing that benefit other than a property right. If, for instance, we hold that laborers deserve financial reward for their work, we could always provide that financial reward in the form of tax breaks, prize money, or government stipends, instead of granting property rights to secure it. Call this the *underdetermination* problem.⁵²

⁵¹ The best discussion of desert-for-labor views is Becker 1993 (Becker ultimately rejects a desert-for-labor view). Hughes (1988) also discusses two desert-based views, which he calls the “avoidance” and “value-added” theories.

⁵² The problem is not always fatal. If it can be shown that labor justifies either property rights or some other option, then property rights are *justified*; its just that labor doesn’t give any special reasons to prefer IPRs over the other option. In that case, underdetermination would be a theoretical infelicity, but not a serious problem. Perhaps there are other reasons, not having to do with labor, to prefer property to other

Existing labor theories of IP deal with these two problems in different ways. One strategy is to give a Lockean style labor-mixing argument for IP. Wendy Gordon (1993), Alfred Yen (1990), and James Child (1997) pursue this strategy. Both Gordon and Yen begin from considerations about ownership of one's labor. Like Lockean labor-mixing views, in intellectual labor a laborer "mixes" their labor with something unowned – in this case, an idea taken from the intellectual commons – which then becomes annexed, through this mixing, to them. In Yen's words, the products of intellectual labor are "practically an extension of the author herself" (Yen 1990, 547). Given this mixing, depriving individuals of property rights over the ideas that are the result of their labor is tantamount to taking their labor: "If you take the objects I have gathered you have also taken my labor, since I have attached my labor to the objects in question" (Gordon 1993, 1545).

Child offers a different labor-mixing view. For Child, intellectual labor is a kind of creation *ex nihilo*, and so gives the laborer strong rights over her products:

An inventor of my acquaintance holds seventeen patents...the most interesting thing about my friend's ideas as property – indeed, I consider it amazing – is that he created them *ex nihilo*. Property was created out of nothing but mental labor. He didn't even need raw materials – as the farmer needs seed, or the potter needs clay. Of course, he needed a pencil, paper, a drafting board, and a slide rule (my friend is an old-fashioned inventor). But these are more like tools. He needed no "stuff" with which to mix his labor. (Child 1997, 67-68)⁵³

entitlements, and these, in conjunction with a labor argument, are sufficient reason for property. I owe this point to an anonymous referee.

⁵³ Such theories are usually referred to as *makers' rights* views. Tully (1983) is considered the locus classicus for such a makers' rights interpretation of Locke's property theory; Sreenivasan (1995) also offers a makers' rights view. For discussion of makers' rights views in general, see Simmons 1998. Rand (1986) offers a similar argument for IPRs, though much less sophisticated than Child's. There are serious questions about the notion of "value" at the heart of Child's view; for critical discussion, see Hettinger 1989: 35-40 (and also Mossoff 2012 for a rebuttal to Hettinger).

Child's argument is that intellectual labor does not remove anything from the stock of stuff that there is in the world. What intellectual labor does is increase the amount of value in the world, without destroying or depleting any existing valuable resources. Intellectual labor takes existing ideas and builds on them to create new works of art, or invent new machines. Because ideas are nonrival, the use of these ideas does not deplete them in any way; the basic scientific principles behind the internal combustion engine, for example, were not depleted when the engine was invented. They can be used again and again. Further, the design for the engine was not depleted when the first engines started rolling out of the factory. It can be improved on, tweaked, and redesigned. In that sense, intellectual labor creates value (practically) *ex nihilo*.⁵⁴

Another view is Adam Moore's Lockean theory (Moore 1997). Moore's view differs from the previous, in that in his account laboring to create some invention or artistic good does not in and of itself generate a right. Rather, laboring creates a "weak presumptive possession and use claim" in a labored-on item (Moore 1997, 84). What is distinctive about labor, for Moore, is that it is one of a group of acts (such as first occupancy) that generates weak, *prima facie* claims to items in the commons. These claims are *ultima facie* justified only if the appropriation satisfies the Lockean proviso to leave "enough and as good" for others, which Moore interprets in terms of Pareto efficiency: "The base-level intuition of a Pareto improvement is what lies behind the

⁵⁴ Of course an obvious objection is that intellectual *labor*, as opposed to intellectual *products*, requires a significant amount of resources. Biomedical research, for example, requires capital, materials, labor, and training. But there is a difference between the use of resources as *investment* and their use as *supply*. Child's argument is an argument about supply, not investment. There may be an objection here, but it would require an argument that investment of resources to support intellectual labor at least dampens the claim laborers have on the results of their work.

notion of the proviso. If no one is harmed by an acquisition and one person is bettered, then the acquisition ought to be permitted” (Moore 1997, 85). Moore treats the proviso as a sufficient condition for justified appropriation.

So the mechanics of appropriation, for Moore, are roughly: by laboring, one gains a weak presumptive claim to an item, which is justified (and thus a right) if appropriation of the item is a Pareto improvement over leaving the item in the commons. In Moore’s paradigm case, an inventor, Ginger, comes up with a new gathering technique that allows her to gather with less labor. The inventing gives Ginger a prima facie claim to the exclusive use of the technique; further, by appropriating the technique, no one is made worse off, because Ginger’s appropriation does not hamper the ability of anyone to invent their own technique (indeed, it might even help them, because whereas Ginger’s technique is proprietary, her inventing process is not, and they can use her strategies of invention and discovery to come up with something of their own). Ginger’s appropriation doesn’t make anyone worse off, and makes her better off, thus it satisfies the Lockean proviso, and so Ginger has gained an IPR over her technique.

I am not going to discuss or criticize any of these views at great length here. The main purpose of this chapter, as I said above, is to present a new labor theory of IP, the productive capacities view. This view is a marked improvement over existing labor theories, as it avoids problems with existing labor views. I will discuss these improvements in section 5, after the new position is on the table.

2. Why Labor Matters

What makes labor a plausible source for a claim to an IPR (or any kind of property right)? At the most abstract level, labor is a distinctive kind of action; it is the use of one's productive capacities in purposive activity.⁵⁵ Labor is different from aimless toil in that, in laboring, the laborer aims to bring about some end; labor is action done with the purpose of bringing about some change in the world. Someone who takes a heap of wood and works with it to make a piece of furniture engages in the activity because they *intend* to bring about a piece of furniture. This is different from someone who tinkers with a heap of wood aimlessly, and happens to create a chair; the laborer *intended* the result, the tinkerer simply *happened* to bring it about. This is not to rule out the possibility of what we can call "productive tinkering", in which the tinkering is meant to yield useful knowledge (as when an inventor tinkers with a toaster to figure out how it works) or try out some strategies or tools, or "brainstorming", where an individual tries out a number of possible ideas, without knowing what the exact end result will be. These acts are purposive in the way discussed here, and so would count as labor as I've defined it, even if the precise result (such as designing a better toaster) was, strictly speaking, unintended.⁵⁶

In laboring, the laborer exercises their capacities intentionally and with forethought. If the tinkerer ends up creating a chair out of a heap of wood, they certainly had the capacity to do so; but the tinkerer did not intentionally use those capacities. The laborer, on the other hand, did; the laborer *chooses* to exercise her capacities, makes a

⁵⁵ This characterization of labor comes from Simmons 1992.

⁵⁶ Thanks to Bill Hasselberger for this point.

plan as to how they will do so and how they will utilize the resources and tools at hand, and *revises* those plans in the light of exigencies that arise during the laboring. To be sure, not all of the process is calculated; at different points during the act of working, the laborer can daydream, or get into a state of “flow”. But the whole process, taken together, has all of the hallmarks of intentional, rational action; there is a purpose, a plan, and a series of choices and decisions.

The capacities at work in labor involve the exercise of skills the laborer possesses; the utilization of tools and resources; and time and energy (more or less of each, depending on the task). Labor also involves opportunity costs, as there is always something else the laborer could be doing (potentially, maybe even usually, more fun or profitable). It also involves at least some training, apprenticeship, and practice, to get and hone those skills, which in turn involves decisions about what to do with time, energy, and resources.

Labor, then, as an activity, is linked with all the ways in which human beings are intentional agents. Labor is itself a paradigm of intentional action, but the productive capacities of an individual are also bound up with a series of decisions over a course of time they made about how to use their time and energy, how to dispose of their resources, and (perhaps) how best to utilize their natural skills and affinities. In the case of labor that is done as part of what we’d call a person’s career, these decisions may have been elements in a Rawlsian rational plan, and the actual laboring – the exercise of the individual’s productive capacities in this particular task – would thus be connected to

some of their most significant life choices, goals, and even identity.⁵⁷ Even at a less grandiose scale, labor involves these sorts of things, as the end of laboring is likely to be tied up with goals the individual has, such as supporting their family, or earning some money to take a trip or buy something they desire.

There is thus an important connection between control over one's labor (as a distinctly human, intentional, rational activity) and the capacity to control and order one's life according to one's own goals, values, and identity. Decisions about how and under what circumstances to use one's productive capacities are among the most significant decisions one can make; laboring, and acquiring productive capacities at all, involves decisions about how to use one's time and energy, what to do with one's life, what one values the most, and what ends are most worthwhile. Think about how much time most people spend at work, the importance of being able to work – ideally at one's chosen profession – in people's lives, and the value that being productive, being able to exercise one's productive capacities in a meaningful and profitable way, adds to their lives. The more control one has over the exercise of one's productive capacities, the more control one has over one's life. The more the exercise of these is free, in accord with one's values, and results in ends that are valuable (either in themselves, or as means to certain benefits), the more one will live a free, valuable, profitable, and even *meaningful* life.

⁵⁷ Rawls, though, may not think property rights over the products of one's labor are necessary to respect these plans. It's debatable whether the historical Rawls (or any Rawlsians) would accept this line of reasoning as a basis for Rawls' recognition of rights to personal property as basic (though here, I don't make a distinction between rights to personal property and rights to productive capital. IP is certainly productive capital, so even if the line of reasoning were acceptable to Rawls and/or Rawlsians, it wouldn't necessarily apply to IP).

This means that, insofar as we value the ability to control and order one's life in this fashion, and value institutions that are conducive to this, we should want institutions that allow individuals the maximum amount of control over the exercise of their productive capacities. Restrictions on the ability of individuals to choose how, when, and under what circumstances to exercise their productive capacities are restrictions on their ability to exercise choice over significant portions of their lives. There is, to be sure, no such thing as absolute freedom here; luck, and natural barriers, will always place some restrictions that can't be removed. But we should want to minimize and ameliorate these restrictions, for doing so is equivalent to creating the conditions in which individuals can control and order significant portions of their lives as they see fit.

This is what is distinctive about labor, what makes it a plausible basis for property. The most valuable feature of property rights is that they help provide the means for individuals to have control over the circumstances in which they exercise their productive capacities.⁵⁸ Suppose I labor on a heap of wood and create something valuable and desirable. If I have a property right over that product, I can control who, and under what conditions, gets access to it. The benefits of my labor – the valuable and desirable product of it – are under my control. This benefits me in two ways. First, I have a

⁵⁸ One could object here that there are some cases where having this kind of control can have the opposite effect, and in fact be harmful to individuals. This goes for a number of rights and liberties; rights to free expression are as much rights to stick one's foot in one's mouth as they are rights to meaningfully engage in political debate. That having these rights also could give individuals too much freedom, and so lead to harms down the line, doesn't mean they are not an important part of institutions which give individuals the liberty to live their lives in accord with their own freely and reflectively chosen ends. Thanks to an anonymous referee for raising this point.

bargaining chip when it comes to setting the terms and conditions under which I exercise my productive capacities. Anyone who wants the valuable and desirable things I can make will need my consent, and so will have to offer me something in return. I now have gained the ability to set the terms under which I labor, which gives me a measure of control over the exercise of my productive capacities, as well as the ability to profit from that exercise. This has an added bonus, of making all the training and practice, and the investment in things such as education, retroactively *worthwhile*. Even if all of these things were not valued in and of themselves, they have now paid off, and so these decisions, investments, and opportunity costs were, retroactively, worth it.

Second, my property right over the end product of my labor now gives me the ability to *exchange*. I can sell or trade the product, in exchange for things that I need and want.⁵⁹ I am able to use my productive capacities to meet my needs, satisfy my desires, and pursue my goals. Through the exercise of my productive capacities, I have gained a measure of ability to live my life as I see fit, by pursuing those ends, whether they are high-toned or just simple pleasures, that I value.⁶⁰

⁵⁹ A similar point, connecting IPRs to the right to contract, is made by Hugh Breakey (2009).

⁶⁰ These considerations hold in different social contexts of labor. Even in what Andrew Reeve (1986) has called the “integrated production with wage labor” model, in which firms own the means of production (in this case, equipment, labs, music studios, computers, and the like) and labor is sold, the ability to use one’s productive capacities to produce something of value (namely, IP) makes the exercise of those capacities valuable, and gives one both bargaining power and the ability to trade the use of one’s capacities for what they value (a certain kind of job or career) or just the means (wages) to get what they value and pursue their ends (saving up enough to set up their own business, say, or travelling or pursuing a hobby).

Another way to put this is that allocating property rights are an important part of respecting individuals as *project pursuers*.⁶¹ Some of the ends that individuals pursue are temporary; some persist across their entire lives (or at least significant chunks of their lives) and structure a range of their actions and decisions about how to use their time and resources. These ends – projects – are of fundamental importance to individuals; they not only are important determiners of how they use time and resources, but also give meaning and purpose to their activities and even shape their identity. We can see this when we note that we can often characterize individuals and explain their actions in reference to their projects. “She is a vegetarian”, “He is a good father”, “She collects African objets d’art”, “He is taking care of his wife’s sister”, explain a great deal about a person; why they did what they did, and what they care about and value. In securing the two conditions described above, property rights allow individuals to use their labor to secure the means to pursue their projects. They can parlay their labor into the means to support their families, buy organic produce, take collecting trips to Africa, and help out relatives in need.⁶²

⁶¹ This is the key plank in Loren Lomasky’s (1987, Ch. 6) defense of property rights (indeed, rights in general). The notion of a project deployed here comes from Lomasky 1987, and is related (through Lomasky’s work) to Williams (1973).

⁶² Though largely similar, the argument for property here differs from that given by Lomasky (1987), in that Lomasky focuses on the importance of having stable and secure access to external objects that are necessary to pursue projects, whereas I focus on the importance of having control over objects for gaining control over the use of, and the ability to benefit from, one’s labor.

3. The Productive Capacities View

This gives us the basis of an argument for IPRs. To see this, consider the following case. Take an individual, Jane, an engineer who has labored intensively to create a new machine. Suppose that we have a full suite of property rights over tangibles, but no IPRs. In such a situation, Jane would have a property right over her machine, but she would have no *intellectual* property right over the design for the machine. If Jane's machine could not be replicated, then Jane's old-fashioned property right over the machine would give her the two benefits discussed in the previous section. Anyone who wanted Jane's machine would have to bargain with her for it, thus giving her the ability to set the terms under which she makes Jane machines, and use them to get in return the means to satisfy her needs and wants.

But of course Jane's machine *can* be replicated. The machine she built realizes Jane's plan for design and construction. Even if she never wrote this plan down, and never told it to anyone, someone with the know-how could reverse engineer Jane's machine and figure out how to build one of their own. But further, anyone can build *and sell* them to others. This includes Jane; Jane could set up a business to make and market her machines under the "Jane's" brand. That brand could have some cachet, given that she's the inventor; and she can even trade on that cachet (assuming there is some restriction on branding) and sell her brand to a bigger, more capable manufacturer (instead of going through the trouble of setting up a business on her own). However, her ability to benefit from her machine is severely limited. In the absence of an IPR over her design, anyone can take Jane's design, make a machine from it, or buy one from a

manufacturer who is not required to give Jane any cut of the profits from sales (or even pay an upfront licensing fee to use her design).

Were Jane's design protected, anyone who wanted to make and sell a Jane-designed machine would have to negotiate with her for the rights to the design, and compensate her accordingly. But without IPRs, anyone (including Jane) can do so, and so Jane is seemingly no better and no worse off than any other putative manufacturer. In fact, though, Jane is worse off. In order to sell her own machine, Jane would have to attract significant capital in order to set up her manufacturing concern. Even if she could do that, it would take quite some time (at least) for her to get her plant up and running, and get Jane machines out for sale. Older, bigger, established manufacturers could simply copy Jane's design and beat her to market. Even if they couldn't get their machines out before Jane's, they could almost certainly offer the machine at a lower price, since they could make a greater number at a lower initial cost; thus, the price of a non-Jane made machine should be lower, due to an economy of scale. If Jane decided against mass production, and simply made machines one at a time, she would very quickly be priced out of the market. Firms are simply more efficient ways to handle some kinds of production; there is no way that Jane could compete against bigger manufacturers for market share, if she's making her machines one at a time and selling them at boutique prices.⁶³ Anyone who wants, and is able, to make a machine can benefit, and potentially benefit *greatly*, from Jane's work, without having to compensate her.⁶⁴

⁶³ This is a distinctly Coasean line of reasoning, about competition between boutique manufacturers and firms (Coase 1937).

⁶⁴ There is a wrinkle here, about co-invention. Suppose two inventors independently arrive at the same set of ideas. Since both invested in the labor, both should be due the

This point can be generalized. Suppose Jane wrote a great mystery novel. Established publishing interests could make and sell various products without paying any of the cost of actually writing the book: they could make digital copies of it coded especially for use with their own e-reader (thus ensuring that those with that device have to pay them for the book); handsome print editions for those who prefer physical books; or they could serialize her novel and put it into magazines they publish.

What this case shows is that, in the absence of IPRs, intellectual laborers lose the bargaining power that comes with having a property right over the end product of their labor. Because novels, songs, designs, and formulas are nonexcludable, without IPRs creators and inventors cannot control access to the products of their work. This means that they cannot use the valuable end results of their labor as bargaining chips to set the terms under which they exercise their productive capacities, and utilize those capacities to meet their needs and desires, and pursue their ends.

Granting IPRs, however, secures these benefits for creators and inventors. Because ideas are nonexcludable, in the absence of IPRs, creators and inventors will have less control over the exercise of their productive capacities than they would have otherwise; either they can't use the value they create as a means to control the circumstances under which they labor and meet their needs and desires, or they have to

property right. Any ground for assigning it to one over the other (for instance, giving it to the first to file a patent application) would be arbitrary, and so the IPR (whoever has it) would seem to be unjustified. This is primarily a question about design of IP institutions; any set of IP rules would have to deal with this and similar cases, and there are more considerations to take into account when constructing these rules than *just* the moral demands on the institutions. It seems the most satisfying solution is simply just for both parties share ownership of the ideas, although this will have the effect of reducing the benefits available to each, among other complications. However, even joint ownership in this case may be impractical. Thanks to Loren Lomasky for pointing this wrinkle out to me.

forgo creating and inventing altogether as a realistic option for the use of their productive capacities. Granting IPRs reverses this situation; it secures for creators and inventors these two benefits. Insofar as we value institutions that give individuals as much ability as possible to order their lives as they see fit, we should want IPRs.

Could we secure this for intellectual laborers without offering IPRs? Suppose instead of IPRs intellectual laborers were offered a package of rewards, such as cash prizes, or government subsidies. This would help creators and inventors use their labor to meet their needs, satisfy desires, and pursue ends, as they could use the goods they were rewarded with for these purposes. However, this would not allow laborers to use their work as a bargaining chip to set the terms under which they labored. Since they had no property right over the product of their work, anyone could use their inventions, novels, or songs as they saw fit, once they were out. Further, since the rewards offered are set by someone other than the laborer, creators and inventors have little control over the terms of the exchange; they either have to accept whatever rewards are proffered, or use their productive capacities in some other way.

Suppose instead of an IPR, creators and inventors were granted a limited package of control rights and rights of attribution, such as inalienable entitlements to decide where and under what circumstances a work is used or displayed and what modifications are made to it, to retrieve it from its current holder or forbid its use, and to attribution as its creator.⁶⁵ This puts laborers in the same position. Because they are inalienable, they cannot use them as bargaining tools. Further, laborers have limited abilities to veto *some*

⁶⁵ These rights of control and attribution are (somewhat misleadingly) called *moral rights* (sometimes *droit d'auteur*), and are a staple of intellectual property law in Europe. For discussion of philosophical issues involving moral rights, see Beitz 2005.

uses of their work, but these rights wouldn't give them the ability to veto the kinds of uses that infringe on their ability to benefit from their work. In short, only an exclusive, alienable right – in other words, a *property* right – will allow creators and inventors control over the circumstances under which they labor, and the ability to utilize their labor in satisfying needs, desires, and pursuing ends.

4. An Improved Labor Theory

In sum, the solution to the distinctiveness and underdetermination problems offered by the productive capacities view is: what is distinctive about labor is that labor involves the use of individuals' productive capacities, which involves decisions (potentially stretching across a significant chunk of an individual's life) about how to utilize time, energy, resources, skills, and what opportunities to forgo, in the pursuit of ends. The more control individuals have over the exercise of these capacities, and the more they are able to benefit from their decisions, the more control they have over significant portions of their lives, and the more they are able to live their lives in accord with their own freely and reflectively chosen ends. Property rights help secure this control. Because of some of the unique features of ideas, only IPRs will suffice for securing this control for intellectual laborers.

In solving the two problems in this way, the productive capacities view is a significant improvement over other labor theories of IP, as it avoids some serious problems with these views.⁶⁶ First, the productive capacities view is an improvement

⁶⁶ I don't want to suggest that the problems I point out here are insurmountable objections to Gordon, Yen, Child, or Moore. Rather, the point is that the productive capacities view is an improved labor theory because it doesn't raise these problems. Perhaps a revised

over labor-mixing views, as the productive capacities view can explain why laboring provides the basis for IP in some contexts, but not in others. For Gordon, Yen, and Child, what makes labor distinctive is that labor establishes some sort of *connection* between the laborer and the produced good, which forms the basis of a claim to that good. Because individuals mixed their already-owned labor in the good (Gordon), or create the good ex nihilo (Child), or because it is an extension of their minds (Yen), laborers have some claim to the good.

This offers a simple solution to both the underdetermination and distinctiveness problems. However, this solution is unsatisfying. There are some situations in which this connection is present, but where labor-mixing doesn't offer a reason for property rights. Suppose I am at a conference, listening to a paper, and an objection to the argument occurs to me. I write it down, and during the Q&A I give my objection.⁶⁷ Here, I have produced something that is the extension of my mind, or the result of my intellectual labor, or an ex nihilo intellectual creation. But mixing my labor doesn't seem to give reason why I necessarily should have property rights in these ideas. If my friend repeats my argument to another interlocutor (provided he tells that person I'm the source), or the individual giving the paper answers my objection in a revised version (provided I get an acknowledgment in the footnote), they don't need to bargain with me, or even get my permission, to use the idea in order to avoid committing theft or trespass. Why did

version of their views, or a "hybrid" theory combining good elements from each, could handle these sorts of problems. However, as they are currently constituted, both the labor-mixing views of Gordon, Yen, and Child, and Moore's Lockean theory, are susceptible to some pretty significant issues, all of which are dealt with by the productive capacities view. I owe this point to an anonymous referee.

⁶⁷ I owe this example, about academic labor, to Loren Lomasky.

the same sorts of activities that establish a connection between laborer and good in one context (like with Child's inventor) not establish a connection here?⁶⁸

We could argue that there are elements of the context that nullify the connection; perhaps the conventions governing academic exchange have it that comments offered at a talk go into the public domain, and by participating in the exchange I've given my tacit consent to this transfer from my mind to the commons. But this just re-raises the question; now we can ask, why does tacit consent to the rules of the paper presentation nullify the connection? Why does that make things that are the extension of my mind not suitable candidates for property, when they are in other contexts? What is special about labor in those contexts that establishes the necessary connection, that isn't special here?

The upshot of this is that the institutional structure in which some intellectual labor takes place makes IPRs superfluous. Academics are able to trade on reputation and research potential, because these are the things that are most valuable to universities, and it is universities that pay academics. But for those who work in fields where it is the potential *commercial* value of the products of intellectual labor that matters, rights that secure credit and control are not enough. Property rights are what secure a bargaining position in a market, and so it is property rights that are necessary when the benefits of intellectual labor are obtained through the marketplace. So, labor-mixing views cannot

⁶⁸ Of course there's a sense in which these are *potential* subjects of property. If I were to develop my objection into a long book, for instance, the fact that the book started out as an objection at a conference doesn't mean I forfeit any property rights I have in it. But while it is still just an objection at a conference, the following argument about institutional context applies; mixing my labor to produce a comment at a paper presentation doesn't give reason why I should have property rights over it. The fact that I could justify a property right if I turned the objection into a book only deepens the mystery. Why does mixing my labor not give me reason for property when all I do is offer a comment at a conference, but *does* give me reason when I write a book based on the comment? Thanks to an anonymous referee for raising this point.

solve the distinctiveness problem. Labor-mixing seems to offer a reason for IPRs in some contexts but not in others (such as at an academic conference), and so cannot be what it is that is distinctive about intellectual labor that makes it the basis for IPRs.

The productive capacities view doesn't raise these issues. Intellectual labor in certain contexts does not give us the basis for IP because IPRs are not necessary to give individuals control over the use of their labor, or the ability to benefit from their labor. The institutional structure in which some intellectual labor takes place makes IPRs superfluous. Property rights are what secure a bargaining position in a market, and so it is property rights that are necessary when the benefits of intellectual labor are obtained through the marketplace. There are plenty of intellectual laborers who are in this position, and so given that, the productive capacities view supports labor as the basis for IPRs (for them).⁶⁹

Second, the productive capacities view is an improvement over Moore's view, because it offers a more robust account of why labor matters in justifying IP. Moore's view cannot establish why Ginger's labor gives her a *prima facie* claim to ideas she labors on. Why does Ginger need *exclusive* use of her technique? If everyone had equal privilege of use of Ginger's technique, Ginger would not be made any worse off. Ginger can still use her gathering technique, and use less labor to gather food, even if everyone else is doing so simultaneously, as her technique, like all potential subjects of IPRs, is

⁶⁹ This may seem to raise the distinctiveness problem again, as it means there is only a connection between labor and IPRs in some cases. But this is only a problem if we think that there has to be some non-contingent, non-contextual connection between labor and IPRs; for example, for any view that held that IPRs were "natural" rights. This is not the view offered here, so is not a problem for the productive capacities view. Thanks to Bill Glod for raising this objection.

nonrival.⁷⁰ If use of her technique precludes use by anyone else, then its clear how Ginger could have a claim to exclusive use of it. But if Ginger's technique can be used by anyone, then there's no reason why her labor should give her any claim to the idea, to the exclusion of others. There's no reason why intellectual labor, then, should give anyone a weak presumptive claim to ideas; Moore's view thus fails to solve the distinctiveness problem, as he doesn't explain why laboring on an idea gives one a claim to it. Further, even if we accept that Ginger *did* have some claim to her ideas, its still not clear why Ginger should have a property right over her technique. Since Ginger does not need exclusive rights in order to use her gathering technique, there's no reason why Ginger needs a property right to use it. Even if we accepted that Ginger did have some claim, it's not clear why this gives us a reason for Ginger to have an IPR.

Again, this isn't a problem for the productive capacities view. What matters is whether or not giving Ginger an IPR will help her gain control over and the ability to benefit from use of her productive capacities. Using her technique to gather food is one way use of her invention helps her do that; selling it to others, using her property rights in it to set up her own business, further refining of the invention to create new products, and contracting her inventing ability to manufacturers of gathering products are important as well. In order to have these Ginger must have an IPR.

This last point leads to an important objection. We could argue that IPRs are not necessary for securing these conditions for individuals, just necessary for securing these conditions for exercise of *intellectual* labor. They could always choose to do something

⁷⁰ Perhaps Ginger would lose her competitive advantage over other gatherers, but the proviso, as interpreted by Moore, should rule out such an advantage as justification for the right, as having an advantage makes other gatherer's worse off.

else with their labor if they can't benefit from creating novels, designs, or formulas. If Jane or Ginger cannot make a living inventing, in other words, they can always just do something else. IPRs don't allow Jane and Ginger to benefit from their labor, they allow them to benefit *by inventing*; but there is no moral requirement to ensure that individuals are able to make a living as inventors.

This is not, however, a very strong objection, for two reasons. First, if we accept that having control over the use of one's labor, and being able use one's labor to further one's own freely chosen ends, are significant moral goals, then it shouldn't matter whether or not creators and inventors have *some* measure of control and *some* ability to use their labor for their own ends in the absence of IPRs. If there are strong reasons pushing against granting IPRs, then the fact that their removal doesn't entail a loss of these two conditions matters; but at this stage no such decisive reasons have been offered, and so the mere fact that creators and inventors will still have *some* of each benefit without IPRs is not a serious objection. Second, there is an assumption behind this objection, which we can call the *assumption of transfer*. The idea is something like: if you are good enough, creative enough, and skilled enough to invent a new machine, write a novel or a song, or the like, then you can use those abilities in an equally productive way on something else. Your skills and abilities transfer easily to other tasks. For some skill sets this may be true; a talented engineer could probably transition into being a talented financier. However, for others, it is not. The ability to write excellent pop songs does not necessarily mean that one can write excellent software code; songwriters are liable to be severely disadvantaged in their ability to use their productive capacities to pursue their ends without the tools to profit from their compositions.

There is another objection here, about the kind of labor involved in some invention and creative work. An important part of what makes something eligible for patent or copyright protection is that it is novel. Many times, such novelty is not the result of extended and intensive labor, but rather incremental tweaks to existing inventions or literary and artistic works. This kind of intellectual labor – marginal innovation – seems markedly different from the kind of exercise of productive capacities by Jane in our example. Marginal innovation doesn't seem to be very much like labor at all; rather, it seems more like *free riding* off another's labor, in that it involves letting someone else do all the hard work of invention, and then adding a small tweak to generate something “novel”. But if we tried to restrict IP only to individuals who produced through “genuine” intellectual labor, and not marginal innovation, we run into all sorts of problems. How could a patent examiner tell the difference between an application that is the result of the former and not the latter? How could we possibly produce evidence of “genuine” intellectual labor? Given the extent to which creative and artistic labor requires building off existing work, could any artist or writer really be a “genuine” laborer? Can we really restrict copyright protection only to wildly inventive, original literary works? Is it really the case that *Ulysses* is worthy of copyright, but not *The Big Sleep*, because Chandler wrote it after Hammett wrote *The Maltese Falcon*, and so it's only a “tweak” to existing ideas?

In short, recognizing IPRs in ideas that are both the product of “one percent inspiration and ninety-nine percent perspiration”, such as Jane's design, and marginal innovation, seems to defeat the purpose of a labor justification for IP. It essentially gives some individuals the ability to exploit others' labor for their own gain by free riding off

their efforts. It also establishes only a contingent connection between labor and IPRs. Some individuals will have patents or copyrights on ideas that they labored intensively to produce; but some individuals will have IPRs over ideas that are only tweaked versions of existing ideas.⁷¹

This is an interesting problem, not just because of the objection it raises, but also because of what it implies about labor and labor theories of property. A big part of the intuitive appeal of labor theories is that labor is, in general, an unpleasant activity. Labor is hard, difficult, usually less preferable to other (leisure) activities, and frequently is outright drudgery. There is a strong intuitive appeal to the idea that individuals who toil should be able to benefit, in the form of gaining property rights over what they labor on. Marginal innovation, however, doesn't seem to fit this description. Because it's not "hard work" in the same way that Jane's inventing is, it doesn't seem worthy of IPRs. Intellectual labor, in general, may seem this way. Writing fiction, painting a picture, and playing music are (at least sometimes) inherently pleasant activities, and hardly seem like the kind of work that is worthy of property protection.

The productive capacities view, however, can deal with this problem. As I stressed in section 4, it's not the drudgery involved in labor that makes it the basis for IPRs, it is that labor involves the use of time, energy, and resources, and decisions about how best to use those (potentially stretching across an individual's entire life). This counts even for marginal innovation – in fact, it counts even for cases of what we can call *miraculous* invention. With Jane, for instance, it could be the case that the design simply came to her one day while she was standing in her kitchen, in a "Eureka!" moment of

⁷¹ Thanks to an anonymous referee, for pointing out this problem.

fortuitous and miraculous discovery. In that case, the investment and opportunity costs involved in coming up with the idea are very low. Even if we grant that this is a plausible story for many or even most of the items that make life comfortable and convenient for us, invention would still not be a costless procedure. Even if the moment of discovery was an instantaneous “Eureka!”, actually turning the insight into a workable design takes a great deal of work. And further, the ability to do that work takes years and years of training and experience to gain the necessary expertise. If Jane invented a useful new machine, even if we grant that the discovery was a miraculous insight or a marginal tweak, Jane still had to do a lot of work to turn the insight into a workable design, and a ton more work preceding even that to become a skilled engineer.⁷² There are legitimate worries here about the *level* of compensation individuals can get from marginal innovation. That, however, is properly a question about *how much* benefit a person should be able to derive from holding an IPR, not whether we should grant an IPR in the first place. In short, even if marginal innovation involves less intensive labor than other inventive or creative and artistic work, on the productive capacities view it still possesses

⁷² There’s a possible science-fiction case here, in which Jane is *not* an engineer, and the fully worked out design just pops into Jane’s head. She then copies it down (without really knowing what it is), patents the design, and makes millions and millions of dollars. The creation was totally effortless, and her reward is truly disproportionate to the costs of invention. If such a case actually occurred, I would have to admit that there is no ground for an IPR over that design, and if most (or even just a fair amount) of invention and innovation actually worked this way, then this would pose a serious problem for the theory given here. But in the real world, as opposed to the philosopher’s fictional universe – where humans arise fully developed from swamp slime, men get thermometers implanted in their skulls by mad scientists, and aliens regularly kidnap people and take them to planets where the molecular structure of the clear liquid coming from the tap is not H₂O – invention is not costless, designs for machines do not implant themselves fully-formed in people’s heads, and engineering degrees actually require serious work. If this means my view only applies to planet Earth in this century, and not in all possible worlds, then so be it.

the qualities that make labor distinctive. It may rankle intuitively, in that it looks like marginal innovators are less deserving of property than the Janes of the world, but what makes labor a basis for property is not that property is deserved benefit for hard work, but that property gives intellectual laborers – even marginal innovators – more control over the exercise of their productive capacities.

5. But are IPRs *Justified*?

The productive capacities view offers part of a justification for *acquisition* of IP; that is, it offers reasons why individuals should be able to acquire ideas as property. Fully justifying IPRs requires weighing these reasons against countervailing arguments against IP; determining whether one set trumps the other, or whether they can be balanced. What the view offered above shows is that IPRs are conducive to securing control certain kinds of individuals – that is, those who have the potential to make a living doing intellectual labor of one kind or another – should have over the use of their productive capacities, and thus substantial and significant portions of their lives. An ideal set of IP institutions is, thus, at least in this one big way, conducive to the goal of creating the economic conditions in which individuals have control over how, and under what circumstances, they labor, and can use their labor as a means to meet their needs and desires, and pursue the ends they find most valuable. Or in short: IP is an institution that is conducive to creating the economic conditions for a more *liberal* political society, in the old-fashioned

Millian sense of a society that allows individuals maximum space to order their lives in accord with their own (freely and reflectively chosen) ends.⁷³

I take this to be a powerful argument for IPRs. But this is, to be sure, a highly idealized picture, and it ignores all the ugly parts involved. A perhaps glaringly obvious rejoinder to all of this is that the picture I have sketched above is so far removed from the reality of IP as to be farcical. In the real world, IP is the special province of giant, multinational media companies and pharmaceutical corporations. IP institutions primarily function to protect their interests, not the interests of small inventors and artists.

There is reason to think that this view of the reality of IP is itself a distortion of reality. Small firms in fact hold a great deal of IP, and IPRs are essential to the (somewhat surprising) industrial trend of the resurgence of the small firm, especially in certain sectors of the economy (such as biotech).⁷⁴ Further, though it is true that IPRs help big companies protect their interests, they are far more valuable to the interests of small firms, with a small number of employees. Such firms have little competitive advantage against big companies when it comes to things like manufacturing, branding, and advertising. IP protections give them a chance to enter markets they would otherwise be muscled out of by bigger companies. This part of the hypothetical story of Jane and her machine is, in fact, fairly accurate. There is also a less direct way in which IPRs protect the interests of individuals: IPRs give individuals bargaining power when it

⁷³ Ala Mill 2008. This has an important side effect: the productive capacities view also serves as an *ideal* for criticizing current IP institutions. If the value of these is that they are conducive to a more *liberal* political society, because of the control they give intellectual laborers to order their lives as they see fit, then this gives us a criterion for evaluating and criticizing current institutions.

⁷⁴ These features of the contemporary reality of IP are discussed in Merges 2011.

comes to setting their compensation packages with big firms. Since individual laborers could always walk with their patents or copyrights, firms are required to offer sufficient compensation in exchange for transferring any IPRs that result from individual work to the firm (Merges 1999a). This is exactly one of the benefits of property rights discussed in the argument given above for IPRs.

To be sure, what I've discussed here is only the upside to IPRs. A full justification of acquisition of ideas would have to deal with the myriad objections raised by critics of IP, in order to show not only that there are reasons why individuals should be able to acquire ideas, but also that acquisition is morally permissible. Here I have offered a defense and clarification of the positive case for IP on the basis of labor. This is a significant result, as labor arguments are rife within the discussion of the moral foundations of IP. James Wilson (2012) has said that one of the two legitimate contributions philosophy can make to larger questions about IP is to provide a systematic account of the normative terrain. The arguments given here have aimed to do just that, with this one part of the territory. If they have provided a clear and normatively robust account of how labor can be the basis for IPRs, then they have succeeded.

Chapter 4

Intellectual Property or Intellectual Monopoly?

In the previous two chapters, we covered two of the most popular and weighty arguments for IP. In Chapter 2, we looked at the incentives argument, and two recent, sophisticated criticisms of it. I argued there that these criticisms fail to recognize that IPRs don't just incentivize intellectual labor, they also incentivize investment in and the production of certain *kinds* of goods, and so in evaluating IP institutions we also have to make arguments about the kinds of goods (and the kind of art, culture, and science) we want and think are valuable. What this shows is that there is no purely technical solution to problems of IP policy. There is no way to determine how IP institutions should be (or whether we should have them) solely by figuring out what the most efficient way to incentivize intellectual labor is; we also have to deal with normative questions about the kinds of goods we should have, and whose interests and preferences IP institutions should serve.

In the last chapter (Chapter 3), we looked at labor arguments, and I offered a reconstruction, which I called the productive capacities view, of the labor argument for IP. However, as I discussed in the final section of the previous chapter, the productive capacities view stops short of a full justification for IPRs. On the productive capacities view, IPRs are not natural rights, but rather are justifiable because they are part of institutions that secure the conditions in which individuals can use their time, energy, and resources to pursue their projects. Though labor is a robust moral basis for IP, we can only consider IPRs justified if these considerations stack up against countervailing reasons against IP institutions. Since these reasons often have to do with harms stemming from granting IPRs, justifying IP involves figuring out whether there are

feasible IP institutions that realize and secure IPRs and either avoid or mitigate these harms.

In these two chapters, the labor and incentives arguments are treated very differently than is usual in IP theory. The incentives argument is usually treated solely from what I called the “supply side” – that is, as an argument that IP is necessary to ensure an adequate supply of intellectual labor. Labor arguments are usually treated as competing justifications for IP, as arguments that IPRs are “natural” rights, and thus as sufficient (if they are successful, that is) to justify IPRs regardless of considerations about incentives.⁷⁵ As I’ve treated them, however, the incentives and labor arguments are complementary. Determining the shape of IP institutions, or whether we should have them at all, requires determining not just what’s needed to incentivize intellectual labor, but also what kinds of goods we want to incentivize production of, and what weight to give to the interests of those individuals effected by IP institutions. My reconstruction of the labor argument gives a reason to give a certain weight to the interests of intellectual laborers in having control over the products of their labor. However, whether these interests are sufficient justification for IPRs requires determining whether there are both feasible and defensible IP institutions. This means, again, looking at both existing IP institutions and feasible alternatives, to determine how they balance benefits to intellectual laborers (and other parties) with harms to others. If the combination of this moral accounting process and some institutional engineering can result in both feasible and defensible IP institutions that afford IPRs to intellectual laborers, then we have a workable theory of IP that includes, as a major part, a justification for IPRs.

⁷⁵ Besides this essay, another exception to this usual treatment is Robert Merges’ pluralist defense of IP (Merges 2011).

In this and the next chapter I will consider two of the strongest arguments against IP. Chapter 5 will consider the argument that, because of the effect they have on access to medicines in poor and developing countries, pharmaceutical patents violate principles of global justice, and are responsible – literally – for millions of deaths from diseases like TB and malaria. The present chapter will consider a much less lurid, but still very serious, objection to IP: the argument that IP is equivalent to extending a monopoly over ideas and markets for ideas, and thus comes with all the baggage, costs, and downside of monopoly. This argument – the *monopoly argument* – is as old as theorizing about IP itself.⁷⁶ Both chapters have a common theme. In both chapters I argue that the problems cited in these arguments are not with IP as such, but with the inclusion of what we can call, following Michele Boldrin and Daniel Levine, certain “downstream licensing privileges” as components of IPRs. The monopoly argument holds that granting such privileges to IPR holders is part and parcel of IP; IP is essentially a monopoly, because IP gives IPR holders downstream licensing privileges. My argument here is that this is not the case. IPRs only accidentally, and not essentially, give their holders a de facto monopoly. There is a feasible conception of IP that does not include, as a substantial component, any downstream licensing privileges, and thus IP is not (by nature) a monopoly.

⁷⁶ Not to be confused with Raustiala and Sprigman’s bogeyman, the “monopoly theory of innovation” (Raustiala and Sprigman 2012).

1. The Monopoly Argument

Equating IP with monopoly is almost as old as IP itself. The idea structured a great deal of the debate over IP in the nineteenth century. Some, like Adam Smith, believed that IP, though regrettable in principle, was a justifiable temporary monopoly:

When a company of merchants undertake, at their own risk and expense, to establish a new trade with some remote and barbarous nation, it may not be unreasonable to incorporate them into a joint-stock company, and to grant them, in case of their success, a monopoly of the trade for a certain number of years. It is the easiest and most natural way in which the state can recompense them for hazarding a dangerous and expensive experiment, of which the public is afterwards to reap the benefit. A temporary monopoly of this kind may be vindicated, upon the same principles upon which a like monopoly of a new machine is granted to its inventors, and that of a new book to its author. (quoted in Boyle 2003, 55)

For Smith, IP was justifiable as a recompense for the risks involved in intellectual labor.

Thomas Babington Macaulay, who anticipated a great deal of the arguments about IP made nowadays in a series of speeches to the British Parliament, thought that granting IPRs was far superior as a system for financing the arts and sciences to patronage, which (he believed) was the only feasible alternative: “I can conceive no system more fatal to the integrity and independence of literary men than one under which they should be taught to look for their daily bread to the favour of ministers and nobles...We have, then, only one resource left. We must betake ourselves to copyright, be the inconveniences of copyright what they may” (quoted in Boyle 2003, 54).

The link between IP and monopoly lay behind a great deal of anti-IP sentiment in the nineteenth century. Liberal distaste for monopoly generally transferred to IP; IP was considered a remnant of mercantilism, another way in which government protected certain favored industries and propped them up by stifling competition. IP was inimical to free trade, its benefits only apparent and not worth the costs. The high tide of anti-IP

sentiment corresponded roughly with the high tide of support for free trade (Machlup and Penrose 1950). The *Economist*, for instance, in 1851, fulminated that:

The privileges granted to inventors by patent laws are prohibitions on other men, and the history of inventions accordingly teems with accounts of trifling improvements patented, that have put a stop, for a long period, to other similar and much greater improvements...The privileges have stifled more inventions than they have promoted, and have caused more brilliant schemes to be put aside than the want of them could ever have induced men to conceal...On all inventors it is especially a prohibition to exercise their faculties; and in proportion as they are more numerous than one, it is an impediment to the general advancement, with which it is the duty of the Legislature not to interfere, and which the claimers of privileges pretend at least to have at heart. (quoted in Machlup and Penrose 1950, 24)⁷⁷

The use of patents was especially odious given the tantalizing possibility of other schemes for inducing and funding intellectual labor. The French government, for example, had in 1839 purchased the patent on the Daguerreotype and allowed anyone to license the patent for no cost, effectively putting it in the public domain (Kremer 1998). This offered a way to reward inventors for “hazarding a dangerous and expensive experiment”, without the inconveniences of a monopoly.

The basic monopoly argument is as follows: IP has two components. First, IPRs give their holders property rights, of a kind, over ideas. To have an IPR is to have exclusive rights of use and access to ideas, along with the power to alienate one’s rights. But IPRs also give their holders rights to control *downstream uses* of those ideas, after the sale of tokens or expressions of them. To have a copyright in a novel, for example, does not just entail having exclusive rights to make and distribute copies of the novel. It also gives one the ability to prevent a range of derivative uses of the novel, such as writing a completely different novel containing similar characters and settings. It is this

⁷⁷ Contemporary defenders of free trade and free markets by and large feel the same; see Kinsella 2008.

second power that makes IPRs equivalent to monopoly privileges. Michele Boldrin and David Levine call this second component of IP the *downstream licensing* component (Boldrin and Levine 2008). The downstream licensing component of IP essentially allows IPR holders to control access to markets for goods that realize or express the ideas they have an IPR over. But it also, in a more abstract way, allows IPR holders to control what others do with their ideas. If others want to write fan fiction stories containing characters from movies I hold copyrights to, or tinker with my protected engine design to see if they can improve it, or engage in political satire by borrowing a logo or image I use in advertising, having an IPR allows me to control whether this does or does not happen, and on what terms.

This is, in and of itself, not *necessarily* a bad thing. We need to be careful to distinguish between two versions of the monopoly argument. On the one hand, we could hold that granting anyone monopoly privileges over downstream uses of ideas is, for a variety of possible reasons, a bad thing. On the other hand, we could hold that there is nothing wrong *in principle* with doing this, and that doing so may even be a necessary evil given certain circumstances, but that *in practice* having this kind of monopoly power tends to lead, for a variety of reasons, to some bad things. Mark Lemley, for example, (Lemley 2005) argues that IP institutions generate significant enforcement costs. IPRs are very difficult and costly to enforce; they require constant monitoring, either by their holder or government, of use of ideas, as well as the costs involved in catching, condemning, and punishing patent infringers and copyright violators. Often, the precise nature of the claims associated with IP are difficult to fix, which leads to a great deal of litigation. Enforcement costs of IP are a direct result of downstream licensing privileges,

and are significant enough (for Lemley) to call into question whether IP is a justifiable monopoly.

Most rehearsals of the monopoly argument are instances of this second kind. Granting monopoly privileges – even over and above the usual problems with monopoly – generates a set of costs, which amount to a net decrease in efficiency over a situation in which there were no IP institutions. Contrast this with the kind of argument given by Seana Shiffrin (Shiffrin 2001). Shiffrin argues that, because of the nature of ideas and intellectual labor, there is a strong presumption against monopolizing ideas, which can only be overcome by stronger reasons in its favor. This is so because ideas, by their nature, are most valuable when they are objects of open and shared contemplation:

...intellectual products often require at least some fairly concurrent, shared (though not necessarily coordinated) use for their full value to be achieved and appreciated. Ideas and their expressions are usually most effective when contemplated by many – when their truths are commonly appreciated and implemented, and their flaws discovered and shared. Indeed, there is a social presumption that ideas and expressions are the objects of open dialogue, exchange, and discussion. Attempts to control, suppress, manipulate, or monopolize ideas and information run counter to the intellectual spirit of open public discussions that promote learning and appreciation for the truth. (Shiffrin 2001, 156)

The nature of Shiffrin's objection is that IPRs, because of their downstream licensing component, forestall the common appreciation and implementation of ideas and expressions, and interfere with their use as objects of dialogue and criticism. Her objection is not to the way the scope of IPRs is interpreted, but rather with the very idea of privatizing ideas in the first place.⁷⁸

⁷⁸ A similar argument is made by Yochai Benkler; see Benkler 1999. Shiffrin's argument is in the same spirit as many of the points made by Lawrence Lessig, the foremost critic of IP; see Lessig 2002.

Assessing both versions of the monopoly argument boils down to two questions. The first question is whether the costs of extending monopoly privileges are worth it. If the examples from Boyle and Lemley of the first version of the monopoly argument work, it is because the enforcement costs of IP and the drawbacks of “enclosure” are not worth the payoff. For the other version of the monopoly argument (such as Shiffrin’s), these costs are *never* worth it, and the question is whether the kinds of reasons Shiffrin, Benkler, Lessig, and others cite are sufficient to show that the downstream licensing privileges of IPRs are never justifiable. This question presents us, essentially, with two options. Either the tradeoffs involved with granting IP are worth it, and extending monopoly control over ideas in the form of IPRs is a necessary evil, or they are not, and IP institutions should be done away with.

However, we also have to ask whether the two components of IP – the package of exclusive rights that come with an IPR, and downstream licensing privileges – are separable. Is having downstream licensing privileges, and thus a monopoly, an *essential* component of IP, or is it rather an unfortunate consequence of the way in which IPRs are interpreted today by courts and legislators? That is, can you have IP without monopoly; can you extend a package of exclusive rights to IPR holders, without allowing anyone to monopolize ideas?

The answer to this second question is crucially important for IP theory and policy. The monopoly argument (both versions) is considered an argument *against* IPRs; if patents and copyrights are essentially a monopoly, then any reasons having an intellectual monopoly is too costly are reasons we shouldn’t have IPRs. There is, however, a third possibility: IP, in and of itself, is not a bad thing (and may well be a *good* thing), but IP

institutions, on the other hand – as they are currently constituted, enforced, and interpreted – give patent and copyright holders de facto monopoly *power* over protected ideas, and in virtue of the scope of current patent and copyright entitlements must be reformed. This is the position taken by James Boyle (Boyle 2003, 2010). Boyle is not against IP per se; however, he holds that innovation and creative work depend on a careful balance between allowing private control of the products of intellectual labor (for Boyle, this is to incentivize that work) and a robust public domain. What concerns Boyle is the expansion of the scope of IP. Boyle argues that the expansion of IP over the last half-century – which he calls the “second enclosure movement” – threatens the viability of the intellectual commons, and thus the possibility of creative and inventive work. For this reason, he advocates for a kind of “environmentalism” for information and culture. The scope of IPRs must be carefully circumscribed, in order to preserve a robust intellectual commons.

The possibility of such a third way – namely, that the costs of monopoly are not worth it, and so we shouldn’t accept these costs as a necessary evil, but that this means we should reform, not abolish, IP – depends on whether or not the two components of IP are separable. If IP is essentially a monopoly privilege, then Boyle’s third way suffers from a kind of conceptual error. Assessing the monopoly argument, then, depends crucially on whether both components of IPRs – property rights in ideas, and downstream licensing privileges – are separable. If the costs of intellectual monopoly are too high, and IP is essentially a monopoly, then we should do away with IPRs. But if IP is *not* essentially a monopoly, then the monopoly argument is fundamentally an issue about design of IP institutions, not the ultimate justification of IPRs. The question would then

become: what do IP institutions have to be like, in order to avoid giving IPR holders monopolies over ideas?

2. Two Problems With the Standard Model of Intellectual Property

In fact, the two components of IP are separable. The central reason for holding that they are inseparable involves two errors about the nature of IPRs and the relation between IP and old-fashioned property rights. Recall the discussion, at the beginning of Chapter 2, about the logic of incentives arguments for IP. The basic idea behind incentives arguments is: ideas are nonrival (adding additional users does not generate any additional costs, beyond the production cost of the good) and nonexcludable (once they have been made available to one person, its difficult to exclude any additional individuals from using them). Because they are both nonrival and nonexcludable, the cost of using them is small, and it's difficult to exclude anyone from using them once they have been made available. The costs of producing ideas, on the other hand – the costs of intellectual labor – are often considerable, and it will be difficult to recover these costs once ideas have been produced and released out into the world. For that reason, there aren't sufficient incentives to do any intellectual labor (or at least, the costs of doing it will tend to keep the amount of labor, and the supply of good ideas, below what is desirable) unless some way is provided to allow intellectual laborers to gain some benefit from their efforts. IPRs are one means of providing benefits and compensation, and thus one means of creating incentives to ensure the desired supply of good ideas. If IPRs are preferable to other means of doing so (such as offering prizes), then this is a reason to allocate IPRs over ideas.

The basic incentives argument for IP is structured by a set of presuppositions, about ideas, property, and IP institutions. Collectively, these presuppositions constitute a way to *model* questions about IP institutions and IP policy. This model – what I will call the *standard model* of IP – is made up of three key elements. First, in the standard model, ideas are *public goods*; they are nonrival, nonexcludable resources, and so, like other public goods (such as national defense, or clean water), it will be very difficult to get all those who benefit from them to internalize a share of the costs of producing them, absent something like IPRs or subsidies for intellectual labor. Second, the essential function of IPRs is to create an *artificial scarcity* in ideas, by allowing certain individuals (intellectual laborers) to exclude others from their use with the aid of legal institutions. Third, questions about IP ultimately turn on the virtues and vices of *private provision* of these public goods. There are other ways of creating the necessary incentives, such as direct government subsidy of research, or technology prize competitions. Whether we should have IPRs depends on whether we should allow privatization of ideas, as a means to secure their supply.⁷⁹ Note that the standard model is not equivalent with (or committed to some version of) instrumentalism about IP. These considerations, in favor of privatization, don't have to turn *solely* on instrumental considerations about the best

⁷⁹ It's worth noting that the public/private split here is, to a high degree, artificial. As Lemley's argument about enforcement costs shows, having IPRs requires a great deal of public expenditure and involvement by government. To have IPRs requires committing public resources in order to facilitate markets in certain kinds of goods, as a way to indirectly finance things like medical research. In a sense, then, creating and financing IP institutions is a kind of government subsidy or public financing of art, science, and culture. This should only pose a serious concern if one thinks that "public" and "private" are binary categories, and not ends on a spectrum. If we think of them in the latter fashion, than there's really no problem calling IPRs a tool for "private" provision of ideas, as using public resources to facilitate markets is far closer to the "private" end of the spectrum than, say, direct government financing, in the form of prizes, for research and development.

means of incentivizing intellectual labor. Considerations about, say, desert or the rights of laborers over their creations may count in favor of IPRs as well, and it may be that arguments have to take into account balancing the interests of laborers with the public benefits of different methods of incentivizing intellectual labor. Regardless of what sorts of reasons are admitted into the calculus, however, the fundamental question remains the same: is the best way (economically, morally, politically, or whatever) of solving the “public goods problem” here the introduction of property rights, and so the privatization of ideas?

IPRs are Rights to Rival and Nonexcludable Goods

There are two problems with the standard model of IP. Both of these problems have serious consequences for assessing the monopoly argument. First, IPRs are not rights to public goods; rather, they are rights to *rival* and nonexcludable goods. This is because, even if ideas themselves are nonrival and nonexcludable, IPRs are not rights to ideas, they are rights to *uses* of ideas. A copyright or a patent does not give the holder rights over the patented or copyrighted idea itself; rather, it gives them rights to, for instance, “reproduce the copyrighted work in copies or phonorecords”.⁸⁰ To have a copyright or a patent is to have exclusive, transferable rights to make and distribute copies or tokens of an invention, certain uses of a logo or other trademark, or manufacture units of a drug. But it does *not* give individuals rights to control the ideas themselves. In a very real way, the actual ideas themselves remain “in the commons”, even though everyone who doesn’t hold an IPR in them is legally prohibited from certain uses of these ideas.

⁸⁰ As stated in the US copyright statute, 17 U.S.C. 106.

This is a plain feature of the way IPRs function. Despite horror stories about, for instance, Disney preventing sales of a comic book with images of Mickey Mouse uttering profanities,⁸¹ IPRs are, compared to rights in physical property, comparatively weak, in terms of control over use and access *to ideas* that they give their holders. A range of uses, even *nefarious* uses, are not prohibited by IP laws, while their physical property-law analogues are. Suppose there is a depressed patent examiner, whose sole ambition in life is to be a great inventor. When he couldn't come up with anything on his own, he joined the patent office, and has spent his life examining the inventions of others. All these years the examiner has kept a notebook, where he's written down the contents of patent applications he admires, hoping for inspiration. One day, he simply snaps. He becomes deluded, and believes the contents of his notebook are *his* inventions. The examiner tells no one; somewhere in his subconscious, he knows he would get found out if he tried to pass off the inventions as his own. But he feels, privately, that he is a great inventor, and is content that his life's ambition has been fulfilled.

If an IPR is a right to an idea, then what the patent examiner has done constitutes theft. He has, quite literally, stolen; he has taken others' ideas, and appropriated them as his own. The fact that no one knows shouldn't matter. If I leave a valuable piece of furniture in my basement and forget that it's there, and a thief breaks into my house and steals it, the fact that I don't know it's gone doesn't mean that it isn't theft. Whether or not there is any potential harm to my interests doesn't matter. If I *never* notice that the

⁸¹ In *Walt Disney Productions v. Air Pirates* 581 F.2d 753 (9th cir. 1978), *cert. denied* 439 U.S. 1132 (1979), Disney sued a group of comic book artists, the Air Pirates, for copyright and trademark infringement, for their use of Mickey and Minnie Mouse in satirical comics in which the characters, among other things, uttered profanities and engaged in an orgy. Disney won the suit. For some discussion of the issues involving copyright and parody, see Waldron 1993.

furniture is gone, then it seems the thief hasn't really harmed me in any way by stealing it, yet the act still counts as theft. But, though what he has done is certainly bizarre, the deluded patent examiner hasn't violated the real inventor's IPRs.

There are, to be sure, ways in which the deluded patent examiner could violate someone's IPR: he could try to make and market one of the patented inventions as one of his own. But in doing so, what he's done is a violation because these are specific uses of ideas that the actual inventors have exclusive rights to in virtue of holding an IPR. To hold that unauthorized and unconsented-to uses of ideas that *aren't* specifically *protected* by an IPR are theft goes beyond what IPRs protect. It is like claiming the toddler who grabs his mother's necklace and yells "Mine!" is a jewel thief; using someone else's ideas to further a delusion, whether innocent (like a child going through a selfish phase) or tragic (like our poor frustrated patent examiner) is not theft, even if it does involve accessing and using protected ideas without permission.

IPRs, then, are not rights to ideas, but rights to uses of ideas. The problem for the standard model here is, whereas ideas themselves may be nonrival, and many uses of ideas (say, contemplating an interesting thought one read in a book, or discussing what happened in last night's episode of *Mad Men*) are nonrival, some uses of ideas are *highly* rivalrous, and these tend to be what is protected by IPRs. Recall the discussion from Chapter 2 about why the "demand" side of incentives arguments matters. Sometimes tokens of ideas are valuable not just because they are tokens of ideas, but because the "packaging" they come in is valuable as well. First edition, signed copies of *Moby-Dick*, the recipe for your favorite dish at your favorite restaurant, jokes told by an excellent comedian – all of these are valuable not just because of the ideas involved, but because

of the packaging they come in. But many tokens of ideas are valuable solely (or at least mainly) because they are tokens of ideas, such as CD's, paperbacks, photographs, newspapers, car engines, and antibiotic pills. There is only finite demand for these tokens; the more of these that are made, the less valuable they become.

Uses of ideas to make these tokens – uses such as making and distributing copies of recorded music – are *not* nonrival. Rather, adding additional users subtracts from their value and adds costs. Having the privilege of making copies of a recording is less valuable if everyone can do it, as the market value of tokens of the recording diminishes as more users have this privilege. Extending privileges for these kinds of uses thus generates externalities, in the same way that overfishing generates externalities; by diminishing the value of the good (that is, the privilege to use the resource in this way), adding additional users effectively “depletes” the good, as future uses will have less value because of the escalating diminishing returns on using the idea.

Thus, whereas ideas themselves are nonrival and nonexcludable, the subjects of IPRs – *uses* of ideas – are rival and nonexcludable. We should, therefore, reject the first two elements in the standard model. The subjects of IPRs are not public goods, and IPRs do not create an artificial scarcity in order to incentivize their production. The goods they protect (certain uses of ideas) are already rivalrous; rather, IPRs just provide a way to give holders exclusive rights to these goods.

Ideas Are Divisible Goods

The second problem with the standard is the way the “privatization” of ideas is conceived. The assumption behind the standard model is that IPRs are like all other property rights, in that they give owners rights to control access to ideas. Property rights are, in Thomas

Merrill and Henry Smith's words, "the simple right of the owner to exclude the world from the resource" (Merrill and Smith 2007, 1853). IPRs, like other property rights, give their holders rights to exclude others from, and control use of and access to, their subjects. Despite the many differences between the two, to have a property right in an idea is, in this one important respect, much like having a property right in a chair or a car. Just as no one can sit in my chair or drive my car without my authorization, no one can do things to my novel or formula or engine design without my permission (subject to the built in limitations on my privileges, such as fair use).

As a consequence of this, downstream licensing privileges are not something extra that gets added on to the enumerated privileges of IPR holders; they are an essential component of having a property right in ideas. Having a property right in a good is not consistent with allowing others use of and access to it, at least not without consent. If someone else can do what they want with my engine design after my patent expires, or lift digital samples from my recordings to make new musical compositions without notifying me, or issue a compulsory license to manufacture and distribute a generic version of a desperately needed drug, then my patent or copyright can't count as property, because it isn't a "simple right" to "exclude the world". Property is *essentially*, in the famous and provocative words of William Blackstone, "that sole and despotic dominion which one man claims and exercises over the external things of the world, in total exclusion of the right of any other individual in the universe".⁸²

⁸² In *Commentaries on the Laws of England* (Blackstone 1765), Book 2 Chapter 1, "Of Property In General".

The problem with this is that ideas, unlike the paradigmatic subjects of property rights, are *divisible* goods. Property rights to material objects have to give owners rights to control use and access to the object itself in order to ensure owners will be able to exercise their privileges in that good. This is because most material objects are not divisible. There's no way I can have certain privileges of use in a parcel of land, for example, unless I can control access to that parcel. Suppose that both of us have common privileges in a parcel of land, and that you only wish to use the land to take a short cut on your way home from work. If both of us had only this one interest in the land, then mutual access to the parcel, and use of it by both of us to take shortcuts home, would not diminish the value of the land to either of us.

In fact, there are many possible uses of the land which do not (necessarily) conflict: having picnics, going for long walks, stargazing, gathering wildflowers, sunbathing, playing football, camping, and staging outdoor concerts, to name a few.⁸³ But there are many uses of the parcel that *do* conflict with these uses: building condominiums, high-intensity agriculture, grazing cattle, and farming Christmas trees, for example. Because of this conflict, it's difficult for anyone to use the land to graze cattle or farm, unless they also have the ability to exclude any potential stargazers and sunbathers. Thus, there's really no way for anyone to have rights to use and access the land, unless they have rights to *control* use and access to the land; that is, to exclude others from doing what they will with the good, even if their demands are fairly

⁸³ This requires that the number of users remains constant. Crowding will occur if the number of users reaches a certain level, even if the demands of the users place relatively low stress on the land. Even if all a set of n users wants with the land is to use it to take shortcuts home from work, there is a value for n that will make the land so crowded that it will be difficult to use it for this purpose, and so diminish the land's utility.

innocuous. The parcel of land is not a divisible good; you cannot give some individual rights to use the land to build condos, graze cattle, and grow soybeans, and allow everyone else in the world to retain common privileges to sunbathe and play football on the same parcel.

But ideas *are* divisible goods. For example, talking about a novel, contemplating its main themes, writing essays about its use of nature imagery, and even writing stories about what happens to the characters after the events in the novel end does not interfere with a copyright holder's privilege to exclusively make and distribute copies of the novel. This is fundamentally different from the parcel of land in our example. Your use of the land to have picnics and take shortcuts is not compatible with my wish to use the land to build condominiums or graze cattle. In order for anyone to have privileges to do the latter, they must be able to restrict others from doing the former. But with ideas, it's the opposite. It is as if *anyone* could sunbathe, have picnics, play football, and take shortcuts without getting in the way of the right holder's desire to build condominiums and grow soybeans. You can, in other words, "peel off" some uses of ideas from others, giving individuals rights to certain uses and letting everyone else retain their common privileges in the others. And in fact, this "peeling off" happens every day, with regard to a range of privileges not specifically protected by patents or copyrights, usually without anyone taking notice. Every time you've watched a movie, contemplated its main themes or discussed its characters with friends, you've used the ideas in the movie without interfering with those uses protected by its copyright and reserved specifically for the holder of the IPR.

We have to be careful here. The natural reaction to this argument is to object that *in reality* IPRs *do* give their holders power to control use of and access to ideas, even downstream uses. Horror stories – like Disney suing to prevent others from using images of Mickey Mouse in satirical comics – do happen, and courts consistently interpret IPRs as giving their holders broad and wide-ranging powers to control downstream uses.

I am *not* arguing that this sort of thing *does not* happen. The above points are solely concerned with the question: are downstream licensing privileges an *essential* component of IP? What the two points I've made above show is that there is reason to believe they may not be. In principle, at least, both upstream (making and distributing copies of novels, for instance) and downstream (writing short stories about what happens to characters after the story in the novel ends) uses of ideas can be peeled apart, and IPRs can give certain individuals exclusive rights to *some* uses of ideas while allowing other (downstream) uses to remain in the commons.

Consequences for the Monopoly Argument

The standard model of IP holds that IPRs provide incentives, and allow intellectual laborers to internalize the benefits of their labor, by providing a mechanism for IPR holders to exclude others from ideas. But, as we've seen, IPRs don't actually function this way; instead, they give IPR holders rights to certain uses of ideas, and allow them to internalize the benefits that come from these uses by excluding others from using ideas in these ways. For this reason, instead of thinking of IPRs as providing incentives by allowing individuals to control use and access to ideas, a better way to think of IPRs is as providing incentives by allowing individuals to extract *value* from ideas, by providing exclusive rights to use these resources in certain (valuable) ways.

We can think of rival uses of ideas, such as making and distributing copies of a novel or drug, as, to use Elinor Ostrom's term, the *resource units* extracted from the underlying ideas (Ostrom 1990).⁸⁴ IPR holders internalize the benefits from the flow of resource units, and the possibility of doing so provides incentives for intellectual labor and the production of goods such as books and drugs. In principle, IP institutions that ensure a flow of resource units to IPR holders (by peeling off certain valuable uses of ideas and giving holders exclusive privileges in these uses), yet allow a set of other uses to remain in the commons could provide a robust set of incentives for intellectual labor, and reward for that labor, to artists, inventors, and novelists. Whether things would actually work out this way is an empirical question, but a set of IP institutions which are strong enough to provide both incentives and rewards yet limited enough to allow robust common privileges is at least *theoretically* possible. In principle, then, the two components of IP can be teased apart; control of downstream licensing is not an essential part of IP institutions that provide robust incentives for intellectual labor.

This is not yet, as it stands, a sufficient reason to hold that the two components of IP are not essentially, but only accidentally, bundled. Even if a set of IP institutions that gives IPR holders exclusive packages of rights to uses of ideas but no (or at least limited) downstream licensing privileges is possible, there may be two further problems. First, such institutions may be so weak, that they don't fulfill any of the functions of IP (such as incentivizing intellectual labor). If this were the case, no functional IP institutions would be possible without robust downstream licensing privileges. Second, it may be impossible to enforce a limited package of rights to uses of ideas without simultaneously

⁸⁴ Ostrom explores possible applications of her theory of property to the ideas and the intellectual commons in Hess and Ostrom 2003.

preventing most downstream uses. So, even if IPRs are limited, exclusive packages of rights to uses of ideas *de jure*, in virtue of practical constraints on enforcement, they might still be *de facto* downstream licensing privileges as well. The upshot of either of these is: even if the two components of IP are not necessarily connected in principle, they are necessarily connected *in practice*. The end result is the same: any functional IP institutions will be (de facto) monopolies over ideas.

3. Digital Sampling, Copyright, and Derivative Works: A Case Study

Let's call these the *watering down* and the *enforcement* problems, respectively. In order to deal with these potential problems, we must consider what kinds of IP institutions would be required in order to create robust incentives for intellectual labor, and give IPR holders some, as Adam Smith put it, "recompense for hazarding a dangerous and expensive experiment". If IP institutions that accomplish both these tasks, and yet *do not* give their holders such extensive downstream licensing privileges that they are (functionally) monopolies, are feasible, then neither of these two problems arises. The bundling of exclusive rights to use ideas with extensive powers to control downstream licensing – intellectual monopoly – is not necessary, and IPRs are not (conceptually) equivalent to monopolies. The key question is: can existing IP institutions be tweaked in such a way that they fulfill these conditions and avoid the two problems laid out at the end of the last section?

To answer this question, let's consider a case in which existing copyright is interpreted as giving its holders strong control over one particular kind of downstream use: sampling of recorded music. The goal is to see if we could remove this part of the

downstream licensing component of copyright while leaving the core – namely, the right to make and distribute copies of recorded music – intact, which avoids the Scylla and Charybdis of the watering down and enforcement problems.

In essence, sampling of recorded music involves lifting a discrete element of a recording and repurposing it, by incorporating it into a new piece of music. Sometimes the resulting works are highly unoriginal; an infamous example is the use of a piece of the David Bowie and Queen 1981 hit “Under Pressure” in the Vanilla Ice song, “Ice Ice Baby”.⁸⁵ Sometimes, however, the resulting works are so novel that the sampled element is unrecognizable. Perhaps the best example of this is the so-called “Amen break”, a seven second drum solo from the forgotten 1970 song “Amen, Brother” by the Washington DC-based funk band The Winstons, which has found its way into hundreds of subsequent recordings.⁸⁶

The aesthetic merits of sampling are still debated, but, for what it’s worth, its difficult to see how sampling differs from a great deal of uncontroversial uses of pre-existing material. Maybe the most vivid example is Brahms’ *Hungarian Dances*. Brahms made use of themes in actual Hungarian folk music in his compositions; a kind of archaic sampling, though in Brahms’ case the “sampler” was his memory. There is little apparent difference between remembering a theme in a folk melody and writing it

⁸⁵ If the reader is unfamiliar with either song, both are available on YouTube. I recommend listening to the David Bowie and Queen song first, as the similarity is apparent after only a few seconds of the Vanilla Ice song; this has the added bonus of sparing the reader from having to listen to more than a short snippet of “Ice Ice Baby”. Note: this fairly blatant lifting of the hook from “Under Pressure” resulted in a nasty copyright infringement suit, which Vanilla Ice and his production team lost.

⁸⁶ For the story, see “Seven Seconds of Fire”, *The Economist* Dec 17, 2011. For the curious, the website Amen Break DB, at <http://amenbreakdb.com/>, maintains a database of tracks containing the famous sample (as well as the original).

down in a new composition, and remembering a drum solo from an obscure funk song and using a digital sampler to lift the seven second break and incorporate it into a new track. Why the use of memory as opposed to a digital sampler should mark the difference between legitimate and illegitimate uses of preexisting works is obscure. Perhaps one could say that Brahms' "sampling" utilized his unique talents as a musician and composer, whereas digital sampling doesn't involve any such faculties or skills. This would require further argument – namely, that having an ear for which samples to lift and how to repurpose them doesn't really involve talent, and that the "uniqueness" of the talents and skills involved in making music contributes significantly to its aesthetic merit. On the face of it, the former seems dubious, and as for the latter there's certainly no necessary connection between talent and aesthetic merit. It's possible for a great musician to use her unique talents in "sampling" folk music, and produce a dreadful composition, while a comparatively untalented producer could use a digital sampler and make a classic.

Digital sampling of recorded music is considered copyright violation, and has been since the infamous *Grand Upright* case.⁸⁷ This is the perfect example of downstream licensing: a copyright holder makes and sells copies of a recording, other

⁸⁷ *Grand Upright Music, Ltd. v. Warner Bros Records, Inc.*, 78 F.Supp. 182 (S.D.N.Y. 1991). Biz Markie, a hip-hop artist, sampled from the song "Alone Again, Naturally" by Gilbert O'Sullivan for a song on his record *I Need a Haircut*. The court granted an injunction against Warner Bros. (Biz Markie's label) on the grounds that sampling constituted "stealing" and was a form of copyright violation. The case changed hip-hop music forever. Whereas previously hip-hop tracks had been built from many samples, the need to obtain permissions and pay copyright holders for each sample used forced producers to build tracks around 1 or 2 samples, instead of sampling freely from a variety of records. The difference is easily appreciated; just listen to a pre-*Grand Upright* hip-hop record, like the Beastie Boys' *Paul's Boutique*, and a post-*Grand Upright* record, like the Notorious B.I.G.'s *Life After Death*. For discussion of the case, and its impact on music, see Boyle 2010, Ch. 6, and Vaidyanathan 2003.

individuals lift discrete elements from that recording and incorporate them into new recordings, the copyright holder uses their copyright to prevent that use and/or extract financial compensation for it. Does this have to be? Can we have a functional version of copyright that allows digital sampling? We need to answer a few questions, in order to figure this out. First, is a piece of music built out of digital samples a “copy” of the original pieces? If so, then there’s no way to give individuals rights to make and distributive copies of music, without forbidding digital sampling. Second, would allowing digital sampling “water down” the incentives (allegedly) generated by copyright? And third, is there any way to enforce copyright without giving copyright holders a de facto power to prevent digital sampling?

It’s tempting to begin with the first question, but this would be a mistake. Doing so assumes that there is some context-free notion of a “copy” that we can rely on, to determine whether pieces of music built out of samples are copies. But any proposed definition of a “copy” or “copying” will be contestable, because it will depend on an aesthetic theory about what the difference is between a *derivative work*, and a mere copy of another work. Derivative works are simply works developed out of or based on pre-existing works, like a screenplay adaptation of a novel or Brahms’ *Hungarian Dances*. A “copy”, however, is more than just “based on” a pre-existing work, but spelling out exactly the difference between a work “based on” and “copied from” another work is extremely difficult.

Consider, as a paradigm case, the relation between Christopher Marlowe’s *The Jew of Malta* and William Shakespeare’s *The Merchant of Venice*. Shakespeare’s play is, without question, in some way derivative of Marlowe’s. But though the plays are very

similar, they are not exact copies of each other. The characters are different, the plots are not exactly similar, the dialogue is different, Shylock is (arguably) a little more morally ambiguous than Barabas. Most critics would probably agree that Shakespeare's play is better than Marlowe's. But in many ways, these plays are not much different. The plot structure is more or less the same, the traits of the characters, the setting (a Mediterranean port city) is the same, among other properties. Indeed, were they to have been written today, it's not hard to imagine Marlowe suing Shakespeare for copyright violation; suits have been brought on much lesser grounds.⁸⁸

If we try to settle this question based solely on our moral intuitions, we find very quickly that it is intractable. This is not just because our intuitions are the result of a confused mix of various and sundry (and often conflicting) norms about intellectual honesty, plagiarism, and copying (though they are that), but also because the question here is, for all practical purposes, unresolvable. Literary works that incorporate similar ideas are not necessarily "copies". It's not enough for a literary work to be similar to another for it to count as a copy and its author a thief; the works have to be similar in the relevant ways. In order to determine whether Shakespeare stole from Marlowe, we

⁸⁸ And Marlowe probably would win the suit, as well, thus preventing Shakespeare from distributing copies of his play or putting on performances of it. Many critics of IP would point to this as a textbook example of how IPRs prevent valuable "transformative" uses of protected ideas (Lawrence Lessig (Lessig 2005), for example, has argued this point vociferously). Shakespeare "stole" a great deal of his material from previous works (this case is not even the most egregious example; portions of *Henry V*, for instance, are directly lifted from Raphael Holinshed's *Chronicles of England*); what better argument against copyright could there be, than the fact that, were he writing today, Shakespeare would have been considered a serial copyright violator, and his works suppressed by legal injunction? This is a point about how copyright is interpreted and enforced, not about what it is by nature. It still remains to be seen if there isn't a suitably modified version of copyright that would have allowed Shakespeare to do his thing (and, for that matter, it remains to be seen whether such a version of copyright is desirable).

would have to determine if the two plays are distinct in the relevant ways, or whether the differences between the plays are (morally and aesthetically) inert. In order to do this, we have to determine which properties are relevant to individuating the plays. In order to individuate expressions, we need a criterion for determining whether or not an expression is, to use Lawrence Becker's term, *singular* (Becker 1993). That is, whether the work is a unique expression of common ideas, distinct from other expressions and possible expressions of those ideas. But a criterion for singularity is, in essence, a complicated aesthetic theory. It has to flow from considerations of what parts of a work are most important to its aesthetic value, as well as theories about originality and authorship.

If the relevant norms are derived from aesthetic views about originality, authorship, and what makes literary works singular, then it's safe to say that the number of answers to the question "Did Shakespeare steal from Marlowe, or merely develop his own unique work from similar themes?" will be roughly equivalent to the number of these theories, and resolving the differences between the answers will depend on determining the "correct" aesthetic theory. At the least, this is practically intractable. But it may just be a fundamentally intractable problem, as it may be the case that any answers to the key moral and aesthetic questions here are essentially contestable. As many critics of IP have pointed out, conceptions of authorship, originality, and singularity, as well as the distinctions between being influenced by a work and outright copying or stealing from another's works, are hopelessly intertwined with various cultural and "folk" aesthetic notions. There may be no right answers here, and the very idea of a right answer may involve nothing other than the reification of one dominant

culture's views.⁸⁹ Even within any particular culture, there are different norms about use of ideas within different communities; for example, the community of scientific researchers may very well have looser norms about derivative uses of others' ideas than the community of visual artists.

However, no such general understanding of a "copy" or "copying" is necessary. If we accept that the point of copyright is to (a) incentivize production of and investment in production of certain kinds of goods (as I argued in Chapter 2), and (b) to give intellectual laborers the ability to use their productive capacities to earn a living (as I argued in Chapter 3), then we can *operationalize* the concept of "copying". If this is the goal of copyright, then a "copy" is anything that, were its production permitted, would dampen incentives and interfere with the ability of intellectual laborers to make a living. The first of our three questions above really reduces to the second; determining whether something is a "copy" means determining whether allowing its production and sale would water down copyright too much. So with regard to sampling, the relevant question is whether producing and selling music built out of samples waters down copyright, to the point that copyright no longer generates any incentives for investment in and production of other kinds of recorded music, and interferes with the ability of musicians to make a living.

What would it take for sampling to have either of these effects? Consider four pieces of music, A, B, C, and D. A is a hard rock song, B is a recording of a jazz standard, and C is a forgotten New Wave track from an obscure mid-80's compilation of

⁸⁹ Siva Vaidyanathan (Vaidyanathan 2003), for instance, argues that hip-hop culture operates with very different norms about borrowing and utilizing others' works, and using copyright to prevent sampling amounts to forcing one community's norms onto another's.

British pop songs. D is a hip-hop track, in which a producer has lifted samples of drums from A, a saxophone riff from B, and some keyboard sounds from C, and combined them with some effects and a vocal track. What would it take for D to dampen the incentives to make and produce A-C, and interfere with the respective musicians, producers, sound engineers, and so forth who made these three pieces of music? It seems the only way would be if D were in some way a substitute for A-C; if someone who would have otherwise paid for A, B, or C, instead bought D. The presence of D on the market for recorded music would have to lower demand for A, B, or C because, for consumers of recorded music, D is substitutable for A-C.⁹⁰

Given this, we can say: D counts as a “copy” of A-C if D is a substitutable good for any of A-C. If D is substitutable for any of A-C, then the presence of D on the market for recorded music will have a negative effect on the demand for A-C, and thus dampen incentives and interfere with the musicians who make A, B, or C-type music to make a living. It’s difficult to imagine this happening. D is a completely different kind of music from any of A-C, and even for consumers who, say, like both jazz standards and hip-hop, pieces of each are not necessarily substitutable (though, they may prefer one over the other, and allocate resources in purchasing recordings accordingly).

This may rankle intuitively. Even if D doesn’t substitute for A-C, there are many for whom sampling still *just seems* like a kind of copying or theft, and it also may just

⁹⁰ Of course, D could also lower demand for A-C because consumers believe D is a superior product to any of A-C, but this can hardly be a reason why D waters down incentives and rewards for making music. The point of copyright is not to keep inferior products on the market and so protect the livelihood of musicians who don’t or can’t make marketable music. If all the “watering down” objection amounts to is that copyright will no longer protect inferior products if we remove downstream licensing privileges, then the watering down effect of separating out downstream licensing privileges from copyright entitlements is nothing to worry about.

seem like there's something wrong with allowing others to make a lot of money by making tracks from samples, while the original musicians who played the instruments get nothing in return. But this is to confuse a number of other issues with the core question here, which is whether or not *copyright law* should allow a copyright holder control over samples from their work. It may be that, for certain communities of musicians at least, sampling counts as copying. This doesn't mean that the norms internal to these communities should be the basis for copyright law. Copyright law should not be concerned with enforcing the norms of specific communities of musicians.

Allowing producers to sample freely without compensating musicians they sample from is a little trickier. It does seem to smack of free riding, in that it looks like the sampler is reaping "the benefits of another's Pains" (to quote Locke) without compensating them for it. But on closer inspection, these intuitions about compensation are misplaced. As Mark Lemley points out (Lemley 2005), no property rights give owners the power to internalize *all* the value generated by their property. A homeowner who plants a beautiful garden cannot sue anyone who enjoys those flowers on their evening walk, because they are not compensated for this enjoyment. When we discussed Lemley's objections to IP at the end of Chapter 2, I argued, contra Lemley, that IP, even in its current, fallen state, does not allow IPR holders rights to internalize the "full social value" of ideas. IPRs give their holders the *chance* to internalize a portion of the value of their ideas *indirectly*, by creating the conditions in which they can extract value from those ideas by producing goods and selling them. If someone buys a record from me, and then uses samples from this record to make a new composition, it's not clear why I'm entitled to anything further. I've already been paid for the copy of the record; why am I

entitled to any further value generated from what others do with the record? This is something like holding that a contractor who builds a home is entitled to a portion of the profit when the owner sells the home, because that profit is in some way due to the labor of the contractor.

But more importantly, even if we granted that there is something “unfair” about it, it still wouldn’t necessarily follow that *copyright law* should correct this unfairness. Again, if copyright law is about incentives and recompense for intellectual labor, then as long as these are unaffected, it’s not clear why copyright should also make sure no “unfair” uses of recorded music is going on. People benefit from positive externalities generated from other’s property all the time. A modest ramshackle home in a nice neighborhood, in which all the other neighbors maintain their lawns and gardens and put additions onto their houses and swing sets in their yards, will be more valuable than the same home in a run-down neighborhood. Is it “unfair” that the owner of the ramshackle home doesn’t compensate her neighbors for the value their hard work in maintaining their homes adds to hers? Possibly. But should the law mandate that she therefore pay royalties on her equity to her neighbors? That’s a separate question entirely, and it would take more than just an argument that there is some “unfair” free riding going on here to establish it.

In sum, it doesn’t look like allowing sampling of recorded music would water down copyright. Music built from samples of pre-existing recorded music is not necessarily a substitute for those pre-existing recordings, and so shouldn’t reduce demand for them. As long as it doesn’t reduce demand for them, a copyright in those recordings is just as valuable as it was before the sampling occurred; having the copyright doesn’t subtract from the value of having exclusive rights to make and distribute copies of that

recording. There may be some sense in which music built from samples is still “copying”, but copyright shouldn’t enforce the aesthetic norms about copying of specific communities of musicians, as (a) the point of copyright is not to enforce said norms, but to provide incentives and the opportunity for artists to make a living, and (b) since these norms are contestable anyway, there’s no way they can provide an objective basis for copyright.

This leaves us with the enforcement problem, but a solution to the enforcement problem here flows naturally from the above considerations. Enforcing copyright is voluntary; a copyright holder has to bring suit against a purported violator, and a court has to agree that the actions of the alleged violator in fact count as a copyright violation. All that would have to happen is to raise the burden of proof. A copyright holder would have to show that the sampler not only lifted discrete elements, but created a work that was, according to our operationalized definition of copying, a “copy” of the work (meaning, similar enough that it was basically the same song, and so was substitutable for the original recording). Copyright holders could thus prevent egregious uses of their work by producers (such as, directly lifting the hook from “Under Pressure” by the producers who made “Ice, Ice Baby”), but wouldn’t be able to sue for copyright violation if someone took a seven second drum solo and packaged it into a new song in such a way that the small sample was unrecognizable in its new context.

4. Reassessing the Monopoly Argument

A great deal of the above, admittedly, is armchair speculation. Whether these proposed alterations to copyright institutions would actually work is entirely an empirical question.

There is no way to tell, from the armchair, whether or not allowing widespread sampling would actually dampen incentives and cut into musicians' abilities to make a living, or whether raising the burden of proof for sampling to count as "copying" would sufficiently deter copyright holders from legally preventing sampling while protecting their rights to make and distribute copies of their recordings. What the arguments here can show is this: in principle, there is no reason why copyright must include, as a substantial component, the power to prevent others from sampling from recordings that the copyright holder makes and distributes. There is a feasible form of copyright that gives its holders exclusive rights to make and distribute copies of music, but not the downstream licensing privilege of controlling who, if anyone, gets to sample from their music, and under what terms. Copyright, then, does not *essentially* include this downstream licensing privilege. It may, currently, give copyright holders this power; but this is a matter of how its *interpreted*, how courts and copyright holders define copying, and what burdens of proof they demand from copyright holders to prove copyright violation has occurred.

IP, therefore, is not essentially a monopoly. The two components of IPRs come apart. There is a feasible conception of IP in which at least some downstream licensing privileges are not included as a substantial part of IPRs, and having an IPR does not give its holder total control over downstream uses of her ideas. If we view IPRs as rights designed to allow their holders to extract value from ideas (by giving them exclusive rights to certain uses of ideas), then granting IPRs is compatible with leaving a host of other uses in the commons, including uses of ideas to produce derivative works for sale (such as new musical compositions built out of samples).

This is *not* however, a total rebuttal to the monopoly argument. What it is, though, is a reason to reassess the monopoly argument, as an objection not to IP as such, but to the way in which IP institutions are currently constituted and interpreted. I have not argued here – in fact, I’ve intentionally avoided – arguing that IP institutions, as they currently are, are not a de facto monopoly. If they are, then the monopoly argument gives us reason to explore IP reform. This is, as I mentioned above, the third way, between accepting IP as a necessary evil or abolishing it entirely, advocated by James Boyle.

My concern here was with the main idea behind the monopoly argument – the idea, going back to the very beginning of theoretical reflection on IP, that IPRs were essentially a form of monopoly privilege. This conception of IP, as intellectual monopoly, should be done away with. The components of IPRs that make them monopolistic – their downstream licensing privileges – are separable, and so there are IPRs that *don't* give their holders substantial monopoly powers. IP is only *accidentally*, not essentially, a monopoly.

Our case study in this chapter focused on copyright, and the power extended to holders of copyright in musical recordings to control whether and under what circumstances anyone sampled from their recordings. In the next chapter, we will take a look at one particular downstream licensing privilege: the ability of pharmaceutical patent holders to exercise control over global markets for generic medicines, extended to them by the Agreement on Trade-Related Aspects of Intellectual Property Rights. I will argue there that this particular downstream licensing privilege should be done away with, in order to bring IP institutions in line with demands of global justice.

Chapter 5

Global Justice, TRIPS, and Access to Medicines

In July of 1861, in the opening months of the US Civil War, the Army of the Potomac marched south from Washington DC, with the intent of mounting an assault on the Confederate capital of Richmond, Virginia. In doing so, they crossed into a region of the United States in which *anopheles* mosquitoes, the mosquitoes that carry malaria, thrived. In the year after Manassas/Bull Run, the first major battle of the Civil War, 1/3rd of the Army of the Potomac suffered from a bout of malaria. Union troops further south were even less fortunate. An expeditionary force that landed at Roanoke Island, North Carolina, suffered a whopping 233% infection rate from the summer of 1863 to the summer of 1864 – which is to say, the average soldier came down with malaria twice. Throughout the Civil War, the infection rate for Union troops, not “seasoned” by bouts with malaria during childhood, hovered around 40%; in one year alone, over 361,968 Union soldiers contracted malaria. Malaria was not a similar problem for the Confederate army; most Confederate soldiers would have contracted the disease as children and, like residents of malarial zones in the world today, either died or acquired immunity to the disease.⁹¹

In the middle of the 19th century, malaria was still a very real presence in the United States. Today, malaria, along with other mosquito born diseases, such as yellow fever and dengue, are virtually nonexistent in the United States.⁹² Many current rapidly developing, middle-income parts of the world are also having progress in combating

⁹¹ This, and the rest of the story of New World malaria, is from Mann 2011.

⁹² In 2010, the US Centers for Disease Control (CDC) received 1,691 reported cases of malaria; 1,688 of these were “imported” (that is, contracted while travelling), 1 was from a blood transfusion, and 2 were “cryptic”. See CDC, “Malaria Surveillance – United States, 2010”. For the story of the eradication of Malaria in the US, see CDC, “Elimination of Malaria in the United States (1947-1951)”. Both can be found at <http://www.cdc.gov/MALARIA/>.

malaria. In Southeast Asia, for example – still very poor but home to some of the world’s fastest growing economies – there were 28 million malaria infections and 38,000 deaths in 2010, a 17% and 15% reduction in the decade from 2000-2010, respectively.⁹³ For the world as a whole, the number of cases per 1,000 (the total incidence of the disease) declined 17%, and the number of total deaths per 100,000 at risk by 26% during that same decade.

For those living in some of the poorest countries in the world, however, malaria remains a source of almost unimaginable misery and suffering. In 2010 there were an estimated 216 million malaria cases in the world, and an estimated 665,000 deaths. The majority of these cases and deaths were in Africa. In that year, for the continent of Africa as a whole, there were 174 million cases and around 596,000 deaths – which is to say, about 80% of the total cases of malaria, and 89% of total deaths from the disease, for that year occurred in Africa. Of those deaths, the vast majority were children: as of today (2013), one child under the age of 15 dies from malaria in Africa every minute.

The bottom 1/5th of the world’s population, those who live on less than US\$1.25 a day, bear a disproportionate share of the total global burden of disease.⁹⁴ Diseases that are treatable or preventable given the resources available in rich countries, such as tuberculosis, malaria, and dysentery, are major causes of misery and death in the world’s poorest countries. Further, diseases such as HIV/AIDS, which are a scourge the world

⁹³ These, and all of the following numbers on malaria, are from WHO, *World Malaria Report 2011*, and WHO, “Fact Sheet: Malaria”, both available at <http://www.who.int/topics/malaria/en/>. Note that the numbers reported here are at the conservative end of the spectrum of possible infections and deaths.

⁹⁴ In 2008, that was 1.289 billion people, roughly 22.4% of the world’s population. See World Bank, “Poverty and Equity Data”, at <http://povertydata.worldbank.org/poverty/home/>. For discussion, see Collier 2008.

over, are a particularly acute problem in poor countries. Simply put, the higher the per capita income of your country, the less likely you or your children are to die from diseases such as malaria, AIDS, tuberculosis, chronic diarrheal disease, and respiratory infections, which together cause millions of deaths a year, almost entirely in the poorest parts of planet Earth. Science fiction fantasies notwithstanding, disease is a permanent part of the human condition, and our susceptibility to it is part of our nature. However, the factors that have the most effect on how big a burden disease is –sanitation, clean drinking water, good public health institutions, and access to medicine and other health resources – are the result of our doings. While the existence of disease is a matter for theological or existential musings, the *distribution* of its burdens raises questions about *justice*.

A great deal of recent work has focused on the role that inequalities in access to medicines, such as antimalarial drugs or antiretroviral drugs for the treatment of HIV, have on inequalities in the global burden of disease. Of special concern is the role of the Agreement on Trade-Related Aspects of Intellectual Property Rights, or TRIPS agreement, which binds all World Trade Organization (WTO) member countries to certain international standards and rules for IP, including pharmaceutical patents. The TRIPS agreement has been blamed for contributing to problems concerning access to medicines in developing countries, and thus exacerbating the burden of disease in the poorer parts of the world.⁹⁵

⁹⁵ For general discussion of issues concerning TRIPS and access to medicines, see 't Hoen 2002; Barton 2004; Bhagwati 2002; Kremer 2002; Pogge 2007; Sell 2001; Shadlen 2007.

In this chapter I consider arguments that the TRIPS agreement, in virtue of its role in contributing to inequalities in access to medicines between developed and developing countries, fails to meet requirements of global justice. The most prominent philosophical work on this topic focuses on the distributive effects of TRIPS. The goal of this chapter is to offer a different argument. My argument is that TRIPS itself, not *just* because of its distributive effects, but in virtue of its internal structure – specifically, the framework it sets for global pharmaceutical markets – is unjust, and therefore that augmenting the status quo through aid projects such as Thomas Pogge’s proposed Health Impact Fund (Pogge 2009) is not sufficient to bring global IP institutions in line with the demands of global justice. Instead, I argue that TRIPS itself either has to go, or must be substantially altered; structuring global pharmaceutical markets by exporting rich-world IP institutions around the world is, from the perspective of what global justice demands, a failure.

In the next section, I briefly discuss TRIPS and the role of TRIPS in contributing to access problems in developing countries. I go on in section 2 to discuss two prominent arguments against TRIPS from its distributive effects, and then give my argument in sections 3-5, before concluding with some brief remarks on the upshot of what is said here for policy discussions about TRIPS. Before I start, a brief word about method. This paper doesn’t present a theory of global justice in trade, or global justice generally, nor does it draw on any general theory. The method here is what Gopal Sreenivasan calls the method of “transitional justice” (Sreenivasan 2007). The goal is to identify some minimal requirements of global justice, which can help set benchmarks for action by policy makers. The argument depends on the viability of the minimal requirements for

global justice in trade I discuss in section 4, not on any general theoretical perspective on global justice.

1. The TRIPS Agreement and Access to Medicines

The Agreement on Trade-Related Aspects of Intellectual Property Rights was part of the broad platform adopted during the Uruguay Round of GATT negotiations, which resulted in the creation of the WTO. The basic idea behind TRIPS is to set standards for IP law and regulation for WTO member states, in order to facilitate international trade in high tech goods, medicines, media, and other products under IP protection.⁹⁶ TRIPS is not the first effort to coordinate patent law across international borders; the Paris convention of 1883 on “industrial property” was the first such international trade agreement, and the forerunner to TRIPS (Helfer and Austin 2011). Prior to TRIPS, however, many developing countries either did not have IP institutions, or had institutions that differed from those of developed countries. India, for instance, allowed patents on manufacturing processes for drugs, but not on pharmaceutical compounds (Barton 2004). Almost immediately, activists and scholars recognized that TRIPS had important consequences for access to medicines and public health in developing countries (Sell 2001). The implementation of TRIPS threatened the global supply of cheap generic drugs, and raised the possibility of huge price increases for medicines in developing countries, with potentially devastating effects on the ability of residents of these countries to gain access to medicines.

⁹⁶ See 't Hoen 2002 and Sell 2001 for a history of the TRIPS agreement.

Why are the effects of TRIPS on access to medicines a subject of moral concern?

There are three reasons in particular that raise concerns about the justice of the TRIPS agreement. The first is the contemporary importance of access to medicines for public health in developing countries. In countries that developed in the 19th and 20th century (such as the United States), improvements in life expectancy were achieved largely through improvements in nutrition, sanitation, water supply and quality, and public health measures. These improvements in turn are a function of development, both economic (as a result of growth in per capita income) and political (as a result of increased affluence of key constituencies that can push, democratically, for improved public services such as sanitation). But similar leaps in life expectancy and health can be achieved today in developing countries through medical technology (Kremer, *Pharmaceuticals and the Developing World* 2002). For example, in 1900, life expectancy for the average American was 47 years. Life expectancy in Vietnam today is 75 years, even though Vietnam has a per capita income smaller than the per capita income of the United States in 1900.⁹⁷ Vietnam's achievement is largely the result of access to medical technology, especially to medicines, such as common antibiotics and childhood vaccinations (Kremer 2002).

Access to medicine, in other words, provides a kind of short cut to ease the burden of disease, which is not (or at least not totally) dependent on broader economic development. That is, access to medicines can greatly help improve the health and well being of individuals in less developed countries, such as Vietnam, even if these countries

⁹⁷ GNI per capita in Vietnam in 2011 was, according to the World Bank, US\$1,270. In the United States in 1900, it was US\$4,096. Data for Vietnam is from World Bank, "Vietnam: Data", at <http://data.worldbank.org/country/vietnam>

lag behind developed countries in factors such as sanitation. This is, in large part, because many of the most pressing health care problems of the developing world are either treatable or preventable. In 2012, the top five causes of death in low-income countries were (in order): lower respiratory infections, chronic diarrheal diseases, HIV/AIDS, heart disease, and malaria. Of these, HIV, malaria, and (many) diarrheal diseases are treatable, as is a common cause of respiratory infections in developing countries, tuberculosis.⁹⁸ It is no wonder that improved access to medicines in countries such as Vietnam can have such an impact on health.

The second is the extent of the HIV/AIDS pandemic in the developing world, especially in Sub-Saharan Africa. Of the roughly 33 million people living with HIV in the world, 22.5 million of them, as of 2009, live in Sub-Saharan Africa.⁹⁹ In 2009 there were approximately 2.6 million new HIV infections in the world as a whole; 1.8 million of those new infections were in Sub-Saharan Africa. That year there was an estimated 1.8 million AIDS related deaths; 1.3 million of those were in Sub-Saharan Africa. For comparison, the number of people living with HIV, new HIV infections, and AIDS-related deaths in South and Southeast Asia, was 4.1 million, 270,000, and 260,000, respectively. Though economic development, and consequent improved public health measures (such as, for example, outreach and treatment for users of intravenous drugs) can have an important impact on transmission of HIV, dealing with HIV/AIDS requires access to anti-retroviral drugs (ARVs, for short). There is simply no way to improve

⁹⁸ Information on causes of death is from WHO, “The Top Ten Causes of Death”, at <http://www.who.int/mediacentre/factsheets/fs310/en/index.html>

⁹⁹ All of the following numbers are from UNAIDS, *Report on the Global AIDS Epidemic 2010: Annex 1*, available at http://www.unaids.org/globalreport/global_report.htm.

quality of life and life-expectancy for those with HIV, not to mention mitigate the economic and political impact of the pandemic, without access to ARVs for those living with the disease. The effects of reduced access to ARVs are, quite literally, measured in millions of deaths. It's safe to say that if one effect of the TRIPS agreement is reduced access to ARVs, then the details of TRIPS are a subject of moral concern.

The third reason is the connection between poverty and disease. Disease, particularly endemic and/or chronic debilitating diseases (such as malaria, AIDS, and tuberculosis) are a major barrier to economic development (Sachs and Malaney 2002). Think of the burden on a single family of caring for a sick member, or dealing with the premature death of a family member. The economic burdens alone are immense: trying to make ends meet given the lost earnings from that member, the cost of food and medicine, and other expenses of caring for sick family members. And this leaves out the emotional and psychological toll of premature death or disability. Now multiply that by 22.5 million, and one can get a general idea of the impact on public health systems, national economies, and civil society in Sub-Saharan Africa of the HIV/AIDS pandemic.

The more susceptible a population is to disease, the less people are in the workforce (due to illness or premature death), the more people need to be taken care of by family members, and the harder it is to marshal large groups of workers together for big projects, such as building bridges or roads. This positive feedback loop generates a situation known as *the disease trap*: poverty exacerbates the burden of endemic disease, which in turn hinders development and contributes to poverty, which exacerbates the

burden of endemic diseases...and on and on (Sachs 2006).¹⁰⁰ Thus, anything that has a negative impact on the burden of disease is at least a contributing factor to the poverty of developing countries. Given the first and second reasons discussed above, access to medicines is a major factor in how burdensome diseases like HIV/AIDS or malaria are to a country, and thus access to medicines is connected in an important way to poverty and economic development. Put simply, there's good reason to believe that lack of access to certain medicines, such as ARVs, contributes to the poverty of certain developing countries, and exacerbates the inequalities in both income and political clout between developed and developing countries.

2. Two Versions of the Access Argument

Call these three issues concerning TRIPS and access to medicines *the access problem*. The access problem is a reason to think there are *concerns* about justice raised by TRIPS, but is not in itself a sufficient reason to believe that TRIPS is *unjust*. Broadly, there are two different perspectives one could take on that question. One possibility is that TRIPS *itself*, as an institution regulating international trade, is unjust. This claim is the focus of a great deal of activism concerning access to medicines, by, for example, NGOs such as Medecins Sans Frontieres (t Hoen 2002). Another possibility is to hold that because of the *distributive effects* of TRIPS, the “status quo” is unjust, and something must be done to remedy these negative distributive effects.

As examples of this latter possibility, consider two arguments for this claim about TRIPS and its distributive effects, one from Thomas Pogge (Pogge 2008) and one from

¹⁰⁰ The disease trap is one of several “poverty traps”. On poverty traps generally, see Collier 2008.

Allen Buchanan, Tony Cole, and Robert Keohane (Buchanan, Cole and Keohane 2011).

Pogge's version is developed from within his views on human rights. Access to medicines is an important part – indeed, given the considerations discussed above, for many people one of the *most* important parts – of the human right to health-care resources. Any institution that raises barriers to access thus violates the human right to health-care resources of those affected. TRIPS, and to a lesser extent pharmaceutical patents generally, raise barriers to access for developing countries, and thus violates the human right to health-care resources of these individuals.

As a remedy, Pogge proposes an alternative patent, which will reward patent holders not by giving them monopoly control over the market for a drug, but by offering them direct financial rewards for impact on the global burden of disease (GBD). Pogge's alternative patent, importantly, will not replace existing patents; rather it will be offered as an alternative, which researchers and pharmaceutical companies can voluntarily accept instead of a traditional patent. A GBD patent requires the patent holder to make their ideas available to any manufacturer of drugs in the world, who can then make and market versions of the drug and sell them around the world for what will presumably be close to the marginal cost of each unit. As the drug makes its way around the world and improves the burden of the disease it treats or prevents, the drug will have a mitigating effect on the global burden of disease. Holders of the GBD patent will be paid in proportion to how big of an impact their drug has on the GBD; put simply, the greater the impact, the greater the reward. Pogge believes this will alter the incentive structure of pharmaceutical R&D in a way that will encourage development of drugs for neglected diseases and cheap pharmaceutical therapies for developing countries, as well as

encourage patent holders to take certain existing drugs off-patent and opt for a GBD patent instead.

To fund this, Pogge proposes the creation of a Health Impact Fund (HIF) to finance the rewards offered to holders of GBD patents (Pogge 2009). The HIF will be financed primarily with transfers from developed countries, though there is also room for contributions from individuals and NGOs. Precise numbers are theoretical at this point, but Pogge and his collaborators Amitava Bannerjee and Aidan Hollis estimate that a functioning HIF will disburse at least US\$6 billion annually from a common reward pool to those with registered GBD patents (Banerjee, Hollis and Pogge 2010).

By contrast, Buchanan, Cole, and Keohane are concerned not with the specific effects of TRIPS as such, but with these as instances of distributive effects falling within the scope of what they call “justice in innovation” (Buchanan, Cole and Keohane 2011). Innovations, such as the development of drugs, are a concern of justice for a number of reasons. They can exacerbate existing inequalities in both wealth and political power, or mitigate these. Further, they can respond to the needs of everyone, or they can be restricted to the needs and wants of a few (mainly, those at the top of the income scale). Both of these dynamics are present in the access problem; IP institutions (such as TRIPS) both exacerbate existing inequalities in access to medicines and in wealth (because of the effects of disease on poverty and development).

Thus, how the products of innovation, such as biomedical research, are distributed, is subject to demands of (distributive) justice. As starting points towards a theory of justice in innovation, Buchanan, Cole, and Keohane propose principles prohibiting distributions of the products of innovation that exacerbate extreme deprivation resulting

from poverty, and what they call “basic economic and political inequalities” – that is, inequalities in income, political power, and access to important forms of social cooperation. The impact of TRIPS on access to medicines fails on both of these counts; it both exacerbates extreme deprivation due to poverty, and existing unjust “basic” economic and political inequalities.

In order to ensure just distribution of the results of innovation, Buchanan, Cole, and Keohane propose the creation of what they call a Global Institute for Justice in Innovation (GIJI). The purpose of the GIJI would be to set rules and policies for the diffusion of innovation, in ways consistent with the demands of justice. One feature of the GIJI would be what they call a “licensing option”: the GIJI would have the power to issue compulsory licenses on innovations, including pharmaceuticals, which are not distributed in ways consonant with the demands of distributive justice. Buchanan, Cole, and Keohane believe that the threat of the “nuclear option” of a compulsory license will incentivize diffusion of innovations without direct action from the GIJI.

Note two features of Pogge’s and Buchanan, Cole, and Keohane’s views. First, neither are direct criticisms of TRIPS *itself*, of the specific internal structure of the agreement, but rather are criticisms of the *distributive effects* of TRIPS. For Pogge and Buchanan, Cole, and Keohane, it’s not that TRIPS itself is unjust; rather, the distributive effects of TRIPS are unjust, and must be remedied. This is not to say that either hold that TRIPS itself is *just*, only that the focus of their argument is on the agreement’s distributive effects, and so both set aside questions about the justice of the terms of the

agreement.¹⁰¹ This opens the door to the possibility that, on both views, TRIPS itself *could be* (even if it currently is not) fully in accord with the demands of global justice, as long as its effects on access to medicines are remedied. Pogge is actually very explicit on this, in a recent statement of his view:

...I believe that it is neither *morally necessary* nor politically realistic to roll back TRIPS in the domain of pharmaceuticals. The preceding arrangements were by no means perfect; and the structural problems of the status quo can be solved through an institutional complement, the Health Impact Fund (HIF), which is specifically designed to resolve failures in pharmaceutical markets. *The crucial moral issue is then not the presence or absence of strong pharmaceutical patent protections, but rather the presence or absence of the Health Impact Fund* (or some similar compensating mechanism). (Pogge 2009, 542) (my emphasis)

The second feature stems directly from the first. Both propose reform schemes that *augment* existing IP institutions and institutions governing trade, but don't require either substantially *altering* these, or abolishing any of their parts. For both versions of the access argument, what global justice requires, as far as TRIPS and access to medicines goes, is a fix for the unjust distributive effects of the agreement.

3. Global Justice and International Trade Agreements

Pogge and Buchanan, Cole, and Keohane's arguments thus approach the TRIPS agreement from a common concern with its distributive effects. For Pogge, these effects are unjust because they reduce access to medicines below what is required by the human

¹⁰¹ Pogge, in fact, argues in a recent article (Pogge 2009) that TRIPS itself may be a good thing; as poorer countries develop and become significant pharmaceutical markets, the existence of robust patent protections in these countries will generate incentives for drugs that treat their particular health needs. Note that this is consistent with holding – as Pogge explicitly does – that the “status quo” (that is, TRIPS minus something like the GBD patent-HIF scheme) is unjust.

right to health-care resources. For Buchanan, Cole, and Keohane, the injustice is a result of the way the effects of TRIPS exacerbate existing unjust inequalities and deprivations.

Both thus imply an answer to the question “What does it mean to say that an international trade agreement, such as TRIPS, is subject to demands of global justice?”. For both, the answer is: it means that agreements such as TRIPS have effects on the distribution of access to things such as medicines, and these effects may be out of line with what is demanded by (for Pogge) the minimal requirements of justice mandated by respect for human rights, or (for Buchanan, Cole, and Keohane) with the requirements of global distributive justice. Both versions of the access argument are thus consonant with a great deal of philosophical work on global justice, which views the distributive effects of economic and political institutions as the main subject of global justice.¹⁰²

The virtues and vices of this theoretical approach aside, I propose to look at the issue from a different angle. In addition to its distributive effects, the TRIPS agreement is also a part of an institution – the WTO – which is itself part of a set of institutions that make up an evolving framework for global governance of trade, politics, finance, and public health. Discussions about global affairs, not just in academic circles but also around kitchen tables and on cable news, are quick to point out the increased interdependence of the world today, over past eras of human history. But our time in history should not just be noted for the increased interdependence of the world, but also for the emergence of a global institutional structure for dealing with planet-level collective action problems. This is not to say that international cooperation is a new thing under the sun; but there are fundamental differences between institutions such as the

¹⁰² The locus classicus for this approach to global justice – as a scaling up of theories of distributive justice to international scope – is Beitz 1999. For criticism, see Nagel 2005.

WTO, WHO, IMF, and United Nations, and past attempts at coordination between nations. The fundamental principle behind the World Bank and WHO, for instance, is that poverty and public health are a concern *of all*; that the economic development of Botswana and the vaccination of children in Pakistan are a concern of the whole world, not just those countries and their allies.

Further, the scope of these institutions is very broad. While it is true that, for example, the WTO is a voluntary organization, given the number of countries involved, and the far-reaching effects on trade of the WTO's rules and agreements, it's safe to say that the existence of the WTO shapes the world economy as a whole, and determines the terms on which global trade is conducted, whether one agreed to those terms or not.

I don't mean to imply here that these institutions are eroding national sovereignty (nor, for that matter, that such a thing would be good or desirable). Rather, my point is that these institutions themselves have an importance beyond just their distributive effects. Their rules set the terms of cooperation on important issues, such as environmental problems and global public health, as well as the terms for actions such as military intervention and global trade. These rules can be fair, inclusive, negotiated by members from positions of equal standing...or they can be the opposite. Whether they meet such conditions *matters*, and not just because of their distributive effects.

When it comes to international trade agreements, these requirements matter for two big reasons. First, whereas it is true that international trade institutions are, as Nagel claims, "a voluntary association or contract among independent parties concerned to advance their common interests" (Nagel 2005, 138), individuals party to the contracting are not the only ones whose interests are affected. In theory, trade agreements are

negotiated by representatives of citizens of their respective countries, but in practice they have frequently been negotiated by individuals who represent only themselves or their class. Sometimes the negotiating parties are the oppressors, rather than the representatives, of their people. To give just one example, both Sani Abacha and Robert Mugabe, two heads of state not often confused with “representatives” deeply concerned to advance their peoples’ common interests, signed on to the TRIPS agreement during the Uruguay Round of GATT negotiations (Pogge 2008). Further, trade agreements are binding not just on currently existing individuals, but also on future individuals as well. Children who did not vote for the politicians who created the WTO nevertheless have to live with the consequences of agreements negotiated before they were born.

This doesn’t mean that the resulting institutions are automatically unjust. But it does mean that trade agreements should, at the least, meet certain requirements that set minimal conditions for them to count as fair and inclusive. We can think of these requirements as (functionally) similar to constitutional rules; they ensure that the terms of trade agreements meet minimum standards of fairness, openness, inclusiveness, and the like, and thus can count as free, fair, and inclusive no matter who is doing the negotiating. In virtue of meeting these requirements, the terms of trade agreements are *endorsable* in principle by those individuals subject to them, regardless of whether they had even an indirect say in their formation.

Second, there are certain norms that are internal to trade, that are operative regardless of whether trade is inter- or intra-national. Trade is more than just exchange; holding someone up at gunpoint and demanding their wallet isn’t trade, even though an exchange (their life for their wallet) has taken place. For two individuals to trade with

each other implies that they enter into the exchange of their own free will, with the intention of furthering their interests, and on terms they believe will be beneficial to both parties.¹⁰³ These conditions – freedom, seeking of mutual advantage, and inclusivity of terms – are, in an important way, *constitutive* of trade. Exchanges that don't meet these conditions, that are coerced, exploitive, or extractive, aren't trade, but are simply extortion, exploitation, and fraud.

Another way to put this is that trade is a *reciprocal* relationship (Lomasky 2007). In trading, individuals part with some good in their control in exchange for something they value more. If each party reciprocates – each giving up something they have, in exchange for something they value more – then the relationship is (ideally) mutually advantageous for both parties. This can be so even if the parties are not friends; in fact, it can be the case if neither party is aware of the other's existence. Mutually advantageous exchange between parties who are not aware (except in the abstract) of each other's existence, in fact, is *exactly what international trade is*, at least in theory. Whether the rules governing the relationship are explicit or tacit, it is because individuals have a reasonable expectation that certain norms governing the relationship will be followed that they submit to parting with goods they control. It is because the advantages of trading *depend* on reciprocity that norms such as fairness and equality of standing are the bases for the general morality of commerce.¹⁰⁴

¹⁰³ The classic statement of this “moralized” view of trade, as not *just* exchange but as *voluntary* exchange for *mutual benefit*, is in Adam Smith's *Wealth of Nations*, Chs. 1-3 (Smith 1981). For a recent discussion of Smith (and other classic political economists') moralized understanding of markets and trade, see Satz 2010, Ch. 2.

¹⁰⁴ There are obvious parallels here with social contract theory; see Lomasky 2007 for discussion of these parallels.

These conditions are not exhaustive, of course. A full discussion of the norms implicit in trade would range far beyond what is appropriate for this paper. This is enough, however, to show that the fact that international trade agreements are negotiated between heads of state and are voluntarily entered into by representatives of states does not mean they are automatically just. Given the far reaching effects of trade agreements on individuals, especially those represented at negotiating tables by despots and dictators, and given the implicit norms governing all trade at any level, international trade agreements have to meet certain conditions that ensure they are free, fair, and set terms that are inclusive and mutually advantageous (in principle, at least). Whatever else global justice requires, it's safe to say that *at the very least*, as a *bare minimum* condition, global justice requires institutions governing international trade that meet certain requirements, such as having inclusive and mutually advantageous terms.

This is similar in many important respects to the view advanced by Leif Wenar, on trade in natural resources (Wenar 2008). Wenar holds that international trade in natural resources has to respect the common property rights citizens of states have to their nation's natural resources. Allowing heads of state to dispose of those resources in ways that enrich themselves and their friends, at the expense of their people, is wrong, regardless of whether or not existing rules permit it. The content of any specific trade agreements notwithstanding, Wenar holds that international trade in natural resources has to meet standards of "clean trade", among which is the requirement that exploitation of a state's natural resources be to the advantage of all citizens of that state.¹⁰⁵

¹⁰⁵ For various reasons, Wenar prefers to think of his view as a theory of "clean trade", not "just trade" (Wenar personal communication).

4. TRIPS and Global Justice

The view on TRIPS here is similar: regardless of the existing content of the agreement or the general terms that frame it, and regardless of its status as a “voluntary” agreement, TRIPS must meet certain requirements in order to meet the demands of global justice. Though of course the distributive effects of TRIPS matter in assessments of it, the *internal structure* of the agreement is also a subject of global justice, independent of its distributive effects. When appraised from this angle, the justice or injustice of TRIPS lies not (just) in its distributive effects, but in the way in which the terms of the agreement structure global pharmaceutical markets.

From this perspective, TRIPS fails to meet minimum requirements of global justice. In essence, TRIPS exports an institutional framework for pharmaceutical markets from developed countries around the world, in the interest of creating a global market in pharmaceuticals. The hallmark of that framework is the ability of a patent to give its holder exclusive rights to a drug, for a limited time, during which they can use their monopoly position to charge high prices for the drug. Setting aside issues about the virtues and vices of patents, structuring global pharmaceutical markets according to the logic of monopoly control over a drug virtually ensures that access to medicines will be a problem in developing and least-developed countries.¹⁰⁶ To see this, consider two effects Western-style patent protection has on markets for pharmaceuticals in developing countries. Pharmaceutical companies usually charge lower prices in developing countries than they do in developed countries. This practice, however, does not (at least not always) redound to the benefit of middle- and low-income residents of developing

¹⁰⁶ Not to mention, among poor residents of developed countries, as well.

countries, but rather to those towards the top of the national income scale (Flynn, Hollis and Palmedo 2009). Given the poverty of individuals who live towards the bottom of the income scale in developing countries, more money can be made by selling *less* units of a drug at a price only wealthier individuals can afford, than selling more units of a drug at a price that a great deal of people could afford (Flynn, Hollis and Palmedo 2009).

This kind of price discrimination is made possible by the lack of any cheap alternative. As long as a patent holder is the only (legal) supplier of a drug, and as long as one can make more money by selling less units and charging x than by selling more units and charging y , they are free to price the drug in a way that restricts access. The only successful way to get around restricted pricing is a steady supply of cheap versions of the drug (Flynn, Hollis and Palmedo 2009). That forces patent holders to reduce their prices to compete, and opens up the market for the drug to more individuals. But as long as patent holders enjoy monopoly position within a market, there is no incentive to drop prices in a way that expands access.

Critics of this line of reasoning point out that TRIPS contains provisions that allow countries to license patented drugs and make their own generic copies, for use *within* their own country, in cases of a “public health” crisis. Further, because the markets for drugs within developing countries are often either very small or nonexistent, patent holders often decline to patent their drugs in those countries. Not only that, but those least developed countries that are members of the WTO are currently exempt from TRIPS, and are not required to become TRIPS-compliant until 2016. The access problem, therefore, cannot be caused by TRIPS and patents, because (a) it exists in countries where TRIPS isn’t in effect, and (b) it exists in these countries for drugs for which there are no

IP protections. The common dominator for countries that have an access problem is poverty, and this is the real culprit when it comes to access to medicines (Attaran 2004, Attaran and Gillespie-White 2001).

This rebuttal obscures, however, an important feature of the way TRIPS affects generic supply. TRIPS allows manufacture of generic versions of patented drugs *for use within particular countries*; it does *not* allow countries to manufacture generics and export them to other countries (Shadlen 2007). Not every country has the technical capacity to make generics; some countries rely on importation of generics made elsewhere, to ensure supply within their own borders of cheap drugs. As long as TRIPS is in effect, and the needed drugs are patented, within those countries that could fill this role of exporter, TRIPS has a chilling effect on global generic supply (Shadlen 2007).¹⁰⁷ And in fact, this is the case, especially with regard to ARVs. “Second line” ARVs¹⁰⁸ are patented in many of the countries, such as India, Brazil, and South Africa, that could make these drugs and export them, thus making it difficult for other countries that lack the technical capacity to gain access to these medicines (Shadlen 2007).

¹⁰⁷ Recently, the International Federation of Pharmaceutical Manufacturers and Associations backed a UN proposal to extend TRIPS extension for least-developed countries indefinitely (IFPMA, “IFPMA Supports Call for Extension of TRIPS Compliance Deadline for Least Developed Countries”, available at http://www.ifpma.org/fileadmin/content/Innovation/IP%20and%20Access/Release_TRIPS_%20extension_10Feb2011.pdf). Given the effects of TRIPS on global generic supply, its no surprise that pharmaceutical companies would support such a measure, as it has no effect on the status quo.

¹⁰⁸ “Second line” ARVs are ARVs taken by HIV-positive individuals undergoing ARV therapy in order to avoid developing resistance. HIV-positive individuals have to take ARVs for the remainder of their lives; developing resistance to “first line” ARVs is thus a serious concern. The introduction of “second line” ARVs into an individual’s management regimen reduces the risk of developing resistance. For individuals who will take ARVs for the rest of their lives to manage HIV, access to “second line” drugs is essential.

What matters, in short, is not the situation involving any particular drug within any particular country, but the structure of the global market for pharmaceuticals. As long as the market for medicines is, at least in large part, such that a single supplier, by virtue of having a patent, has downstream licensing privileges over manufacturing of generic drugs across the globe, and thus can maintain a monopoly position in world pharmaceutical markets, prices and the availability of cheap generic versions of needed drugs will create an access problem for poor residents of developing countries. This is not a side effect or a byproduct of the TRIPS agreement; rather, this is a *direct result* of the agreement itself. By exporting rich-world IP institutions around the globe, TRIPS structures global markets for pharmaceuticals along the same lines as national markets in developed countries. The whole *point* of this framework is to create markets where, for a time, prices can be kept artificially high, by granting market share to a single supplier and forbidding production of cheap(er) alternatives. And it is these two structural features of pharmaceutical markets defined by patent logic that cause the access problem.

Given this, it is difficult to see how the terms of the TRIPS agreement could be considered fair and inclusive. Under TRIPS, global markets in pharmaceuticals are structured in such a way that access to medicines remains a perpetual problem for huge portions of the world's population. When we consider the importance of sufficient access – that is, those three conditions discussed in section 1 that generate an access problem in the first place – it is difficult to see how we could consider the terms of TRIPS endorsable in principle by all those who are subject to it. All evidence points in the opposite direction; the terms of TRIPS are profoundly *unfair* to residents of developing countries, and discount the health needs of poor citizens of these countries, in favor of the

economic interests of pharmaceutical companies and their investors. As the economist Jagdish Bhagwati puts it, "...the WTO ought to be about lowering trade barriers and tackling market access problems. The inclusion of IPP [intellectual property protection] has turned the organisation into a royalty collection agency" (Bhagwati 2002).

The distributive effects of TRIPS matter here, but in a different way from how they matter in Pogge and Buchanan, Cole, and Keohane's arguments. The argument here is not that TRIPS and the circumstances surrounding it are unjust *because* of the distributive effects of the agreement. Rather, the argument is that the distributive effects of TRIPS are *evidence* of the more fundamental injustice of the agreement, which is the way in which it structures global pharmaceutical markets. For these reasons, and not just in virtue of its distributive effects, TRIPS is unjust.

5. Reform vs. Aid

If TRIPS is unjust not only in virtue of its distributive effects, but in virtue of its internal structure, then remedies which correct the negative distributive effects of the agreement, yet leave its internal structure largely intact, are not enough. The upshot of the argument given in section 4 is that TRIPS is simply not acceptable as an institutional framework for markets in pharmaceuticals. This is, admittedly, problematic. As Pogge has stressed over and over, getting rid of the TRIPS agreement raises concerns about political feasibility (Pogge 2009). These feasibility concerns, however, might be a bit overblown. Objections to TRIPS are not new; the agreement has been the source of a great deal of controversy almost since its birth, and was a major sticking point at the last (failed) round of WTO negotiations ('t Hoen 2002). Opposition to TRIPS by developing countries,

NGOs, and activists has been successful in the past; for example, it spurred the Doha Declaration on TRIPS and public health, which “clarified” the WTO’s position that developing countries could use their own judgment as to use of some the public health emergency provisions of TRIPS (such as issuing compulsory licenses for patented drugs).¹⁰⁹ Given past (albeit limited) successes on this front, it’s not clear why pushing for further alteration of the structure of TRIPS isn’t feasible.

In addition, the kinds of remedies for the distributive effects of TRIPS offered by Pogge and Buchanan, Cole, and Keohane raise feasibility worries of their own. Both schemes depend on massive aid projects and large transfers of wealth from currently rich countries to finance them. Given the track record of developed countries on foreign aid, its difficult to imagine that these transfers will be forthcoming any time in the near future. Further, the operations of both have the potential to generate a great deal of rent seeking. Giving a small cabal of politicians and technocrats power to decide what counts as a rewardable reduction in GBD, or power to threaten or impose a compulsory license on a firm, creates incentives for firms to lobby heavily, and do what they can to influence the decision-making processes of both bodies (the HIF and GIJI). There is no guarantee, of course, that this will happen, but given what we know about the workings of actual

¹⁰⁹ The Declaration on TRIPS and Public Health was adopted in the November 2001 ministerial meetings of the WTO in Doha, Qatar. The declaration reaffirms the flexibilities built into the original agreement, such as the right of states to issue a compulsory license for a drug in the event of a “public health emergency”, but does little to alleviate the broader worries about generic supply and global pharmaceutical markets (see Shadlen 2007 on this point). Nevertheless, the Declaration at least makes concern for the effects of patents on health in poor and developing countries a part of WTO policy. Article 4 of the Declaration states categorically: “...we affirm that the Agreement [that is, TRIPS] can and should be interpreted and implemented in a manner supportive of WTO members' right to protect public health and, in particular, to promote access to medicines for all.” The full text of the Doha Declaration can be found at http://www.wto.org/english/thewto_e/minist_e/min01_e/mindecl_trips_e.htm

bureaucracies that have similar powers to affect the long-term financial prospects of big businesses, we can make a reasonable guess that the potential for this sort of rent seeking behavior is a serious concern.

Proposals like Pogge's carry an implicit requirement that the global economic and political landscapes stay relatively the same as they are now. But, as Hume so eloquently observed, "The face of the earth is continually changing, by the encrease of small kingdoms into great empires, by the dissolution of great empires into smaller kingdoms, by the planting of colonies, by the migration of tribes" (Hume 1987, 481). Sub-Saharan Africa, for example – usually the go-to example of desperate poverty – was one of the fastest growing regions of the world in 2011, posting a 5% growth rate despite the troubles for the world economy at large, and preliminary projections for its 2012 growth rate are for growth to have hit 5.5%, higher than any other part of the world except developing Asia.¹¹⁰ Patent reform that depends too much on the continued affluence of currently rich countries, and their continued ability to support massive public aid projects, is not necessarily the best strategy. It is, in a way, a manifestation of the general penchant for aid as a solution to what ails developing countries, instead of supporting structural reforms to the world economy that would, for instance, increase access to world markets for their exports (Easterly 2007).

But more importantly, both offer what is essentially a short-term solution to a structural problem. The problem, as I've argued above, is not just the distributive effects of TRIPS; it is that TRIPS structures global markets in pharmaceuticals in a way that makes it very difficult for developing countries to gain access to medicines. The ultimate

¹¹⁰ See International Monetary Fund, "Regional Economic Outlook: Sub-Saharan Africa", available at <http://www.imf.org/external/pubs/ft/reo/2012/afr/eng/sreo0412.htm>.

goal of action should be not *just* to alleviate short-term suffering caused by disease in poor countries, as admirable as that goal is; it should be to create global IP institutions (and global institutions generally) that don't *contribute* to the burden of disease in the first place, by, among other things, restricting access to medicines. What poor and developing countries need is not reform schemes that make their ability to access medicines depend on the goodness of rich world politicians' and celebrity donors' hearts, but global IP institutions that treat their health care needs and interests as equally important, and structure biomedical research and pharmaceutical markets in a way that enables them to meet those needs without handouts. These countries are not "basket cases" that need "help"; they are partners in global governance and the world economy and need global institutions that are free, fair, and in which their needs and interests are represented.

6. Conclusion

Perhaps the best way forward is to admit that the goal of creating a global market in pharmaceuticals by exporting rich-world IP institutions is a failure. As the lawyer and activist Ellen 't Hoen puts it, "The very fact that public health and access to medicines have been singled out as major issues needing special attention in TRIPS implementation indicates that health care and health care products need to be treated differently from other products" ('t Hoen 2002, 46). TRIPS, after all, also creates barriers to access for DVDs and smartphones, but neither generates an access problem of the kind detailed in section 1.

Rather than furthering the agenda realized in the TRIPS agreement, it's worth considering a threefold plan of action: first, short-term strategies to address shortfalls in access to medicines in poor countries; second, as 't Hoen suggests, giving developing countries broad discretion when it comes to determining whether and on what terms to implement TRIPS requirements for pharmaceuticals ('t Hoen 2002); and third, giving developing countries additional latitude to experiment with homegrown IP institutions, which are more attuned to their particular needs and can make use of local knowledge about public health.¹¹¹ Brazil's extremely successful AIDS program, which does all three, could be considered a model for this approach.¹¹² However, in order to properly tackle such policy questions, we need clearer views of what global justice demands of IP institutions, and institutions generally. Helping get clear on this is the best way philosophy can make a contribution to just to theorizing about global justice, but to *global justice*, full stop.

¹¹¹ Such institutional experimentation has many advantages; for one thing, it allows these countries to make use of local knowledge (such as knowledge about which drugs are needed the most for which diseases and sectors of the population), and adjust their IP laws accordingly. Institutions that can make use of such knowledge generally have an advantage over institutions that are imported and then gerrymandered to fit local conditions; see Boettke, Coyne and Leeson 2008 for discussion. For more on institutional experimentation, see Coyne 2013.

¹¹² On Brazil's AIDS program, see Shadlen 2007.

Chapter 6

Some Conclusions

Though Chapters 2-5 are, again, largely discrete, and can be read independently of each other, with the benefit of hindsight we can see there are clearly some common themes that run throughout. The first and foremost is a general opposition to the approach I've referred to throughout as instrumentalism about IP. Instrumentalism is the view that the central theoretical and policy questions about IP are empirical questions about the instrumental value of IP institutions, as tools for incentivizing intellectual labor. IP is "good" if it is necessary to incentivize intellectual labor, and thus ensure an adequate supply of good ideas, and "bad" if it is not necessary. IP policy thus boils down to a technical, engineering question about the proper design of IP institutions, where what would count as "functioning" or "successful" IP institutions is defined in terms of supply of good ideas.

I directly addressed this view in Chapter 2, but the argument against it is developed throughout. The criterion of "successful" or "functioning" IP institutions, which is invoked implicitly in assessments of the incentivizing goal of IP, is loaded with normative assumptions and buried commitments. Among those explored in this dissertation are: what kinds of goods we want to incentivize the production of; whose interests IP institutions should serve, and whose interests to discount in their design; what norms about copying, influence, and derivative art are relevant to getting functioning copyright institutions; how much the general considerations about economic liberty and liberalism raised in Chapter 3 matter in design and justification of IP institutions; and whether there are some interests or needs – such as the interests most people presumably have in not dying from AIDS or malaria – that trump all others in determining what IP institutions should be like (or whether we should have them at all).

That this is the case may seem obvious to many readers, or at least like a very easy point to make. Its important, for those with extensive philosophical training, who are primed to see normative issues everywhere and behind everything, to keep in mind that it is far from obvious *to others* that this is the case. These sorts of normative considerations are certainly *not* front and center in legal and economic theory of IP. The various pieces of that literature I've cited throughout should be sufficient evidence that this is the case.

But further, the goal here was not just to point out that this is the case, but was also to demonstrate a means for incorporating discussion of these sorts of normative issues into IP theory and policy. Ethical considerations are usually thought to enter into discussion of IP in the form of an either/or: *either* IP institutions are justifiable because they are needed to incentivize an adequate *supply* of good ideas, *or* there are independent moral reasons for IP, sufficient to justify IPRs regardless of considerations about incentives.

The approach to the question developed in this essay is different. Rather than treat the incentives argument, and “moral” arguments (such as labor arguments) as opposed, independently sufficient lines of justification, I treat them as complementary. The anti-instrumentalist approach taken in this essay expands the incentives argument from two sides. First, it expands the incentives argument from the “demand” side; IPRs don't just incentivize the supply of intellectual labor, they also incentivize the production of certain kinds of goods over others. We can't settle the incentives argument without settling questions about what kinds of goods we want, and thus whose preferences to favor and whose to discount. Second, I've expanded the incentives argument to take into

account not just of the interests of “society” in having an adequate supply of good ideas, but also the communities of artists, inventors, and researchers affected by the structure of IP institutions. IP institutions affect how these individuals are able to ply their craft and earn a living, and (as the discussion about sampling in Chapter 4 showed) enshrine certain norms about the making of art, science, medicine, and technology. By incentivizing the production of certain kinds of goods over others, IP institutions favor certain kinds of intellectual labor, and thus certain kinds of *laborers*, over others, and have a deep effect on how creators and inventors structure their labor and ultimately live their lives. This introduces a significant dimension into discussions about what the proper incentives are. The kinds of incentives generated by IP institutions not only affect the kinds of goods that are produced, but also affect the interests of producers of those goods, and so those interests (as well as the interests of consumers of art and medicine) must be taken into account.

In short, assessing IP institutions, and determining whether we should have them (and if so, what they should be like) is multidimensional. It requires investigation along three separate axes: (1) What kinds of goods do we want? (2) How do we settle conflicts of interests and preferences, among the different groups of individuals affected by IP institutions? and (3) What other principles, such as demands of distributive justice, constrain the design of IP institutions? Rather than reducing IP theory and policy to one question, about incentives and the instrumental value of IPRs, these axes are equally important.

The essay is officially agnostic about whether we should or shouldn't have IP; I haven't argued (in fact I've stubbornly avoided stating) that the arguments I make here

are sufficient justification for IPRs. But it's clear that throughout there is definitely a pro-IP slant. In part the motivation for this is dialectical. Even as courts and business have been riding a high tide of pro-IP sentiment, in IP theory and policy there is a great deal – maybe even, at this point, a preponderance – of anti-IP views, a big sampling of which have been cited and discussed throughout this essay. I am firmly in agreement with Hume, that sometimes the best thing you can do in philosophy is point out the merits of both sides of an intense ideological divide, in an effort to moderate the passions each position excites and thus facilitate compromise.¹¹³

For sure, though, one of the conclusions of the dissertation is a cautious endorsement of IP institutions that grant intellectual laborers a limited package of exclusive rights to uses of ideas, but at the same time maintain robust common privileges of use and a certain flexibility in adjusting the entitlements included in the IPR package in response to other demands, such as demands of global justice. This is only, however, a pro-IP position if you ignore the huge gaps between this conception of IP and actual IP institutions as they currently exist and are interpreted. When you consider this gap, and consider how far what we have is from what I present here as defensible IP institutions, the dissertation as a whole is as much a basis for a deep *criticism* of IP as it is a defense of it. Chapters 4 and 5 bring this front and center.

Another major theme of the dissertation as a whole is a criticism of the tendency to look at IP as fundamentally a matter of assessing the proper tradeoffs between net costs and benefits of IPRs. This way of approaching the problem is enshrined in the incentives and monopoly arguments, and lurks in the background of IP theory and policy

¹¹³ See *Of the Coalition of Parties* (in Hume 1987).

generally. A major consequence of the discussion in Chapter 4 is that, once we jettison some of the confusions that come with falsely analogizing IPRs to old-fashioned property rights, we can see that IP institutions are highly malleable. As long as the rights in the IPR bundle allow individuals to have enough control over certain uses of their ideas, sufficient to ensure they can extract a certain amount of value from using ideas, we have IP institutions that can accomplish the twin task of incentivizing labor and allowing intellectual laborers to make a living off their work, and thus use their productive capacities to satisfy needs, desires, and pursue projects.

This allows a great deal of leeway to experiment with IP institutions, to mix and match entitlements in different bundles, weaken or strengthen certain privileges, and either restrict or expand the protectable subject matter of IPRs. Therefore, rather than seeing questions about IPRs as involving tradeoffs between the evils of monopoly or restricting free expression or communication or the like, these are really questions about getting institutions that accomplish what we want, and yet are defensible (in that, they minimize or mitigate harms). Instead of either accepting or rejecting the costs of IPRs, the framework given in Chapter 4 opens conceptual space for seeing costs and goals as reasons to tinker and experiment with IP institutions.

I want to close with a case that perfectly illustrates the central themes of the dissertation. Chapter 5 only stuck a toe into the turbulent waters of debates about pharmaceutical patents. That chapter dealt only with one part of the alleged difficulty with these patents, the access problem. There is another problem, which is part and parcel of this as well: the *incentives problem*. Medicines are very, very expensive to develop, and drug

development is a risky business; many drugs never make it through the various stages of testing and approval. Patents on drugs allow drug companies to charge high prices and bring in huge margins; but (the argument goes) if they couldn't produce these margins, drug companies couldn't attract investment (because of the risks of drug development), and couldn't fund R&D. Without patents, it's hard to see how drug development could be profitable, or even cost-effective; patents are thus necessary (again, so the story goes) for providing the incentives to attract capital to pharmaceutical companies, and fund drug development.

This line of reasoning is so persuasive that even many IP skeptics in law and economics accept that the pharmaceutical industry is a special case, and that pharmaceutical patents may well be necessary (even if other IPRs, like copyright, are not).¹¹⁴ Regardless, the incentive structure created by patents has serious consequences for the availability of medicines in poor countries. Since these countries cannot afford medicines charged at patent prices, there is no effective demand for drugs in these countries; in other words, there is no market there for drugs. Thus, diseases that primarily affect residents of poor countries are not targets of R&D by pharmaceutical companies; there is no money to be made by developing drugs for these diseases, so patents provide no incentive to do so. This is the incentives problem.

The incentives problem is usually cast as a problem about developing new drugs for diseases that afflict the poor. But equally as important, it's a problem about developing drugs for diseases that afflict both rich *and* poor, which are affordable and effective for *both* rich and poor. There is no more vivid and tragic illustration of this

¹¹⁴ See, for instance, Raustiala and Sprigman 2012.

aspect of the incentives problem than the controversy over the so-called “short course” AZT therapy for prevention of perinatal HIV infection.¹¹⁵ Early in its therapeutic history, it was recognized that AZT can help prevent perinatal and neonatal HIV infection – that is, infection of a child by the child’s mother prior to, during, or shortly after birth. An intensive therapeutic regimen involving AZT, nicknamed the “long course”, was determined to be effective for this purpose, and prescribed as a prophylactic against perinatal and neonatal infection beginning in the mid-1990s.

However, at the time, AZT was still under patent by the pharmaceutical company Burroughs-Wellcome (now GlaxoSmithKline), and the cost of the drug was astronomical. Because of how much AZT was needed for the “long course”, and because the cost of AZT made it highly unlikely that the “long course” could be used for prevention of perinatal and neonatal HIV transmission in poor and developing countries, the CDC conducted clinical trials of a “short course” AZT regimen, beginning in Thailand in 1994. The “short course” trial sparked international outrage in the medical community, and among bioethicists and public health officials. The reasons were myriad; primarily, the outrage centered around the charge that clinical researchers were putting the lives of infants at risk by administering a placebo as part of the research studies. But one big source of outrage was over, for lack of a better description, the perceived injustice of the whole situation: an effective treatment for a serious and tragic health problem (that is, transmission of HIV from HIV-positive mothers to their infant children) had been discovered, the cost of the treatment was beyond the means of the vast majority of

¹¹⁵ The bioethics literature on the “short course” trials is huge. The story here is aggregated from: Angell 1997; Annas and Grodin 1998; Crouch and Arras 1998; and Lurie and Wolfe 1997.

women who have HIV (who then, as now, live in poor and developing countries), and so this great leap forward in HIV treatment would not benefit the great majority of people who have HIV. Instead, the CDC used public funds to try and develop a more cost-effective therapy – solely because it was more cost-effective and sustainable, not because it was a therapeutic improvement. The poverty of countries where the majority of HIV-positive women live meant there were no incentives for developing a treatment for mother-child HIV transmission for use in these countries; though it was a problem for both rich and poor, the incentives existed only for developing a therapy that benefitted the rich.

The incentives problem for drug development encapsulates all the different dimensions of assessing IP institutions discussed in this essay. It's clear from this case that we cannot look at pharmaceutical patents and assess whether they are necessary or not for incentivizing drug development and financing of pharmaceutical R&D without also asking “incentives *for which drugs, for which diseases, and for whom?*”. Patents could be excellent instruments for incentivizing drug development, but fail on every relevant normative metric because they incentivize the wrong kinds of drugs, for the wrong kinds of diseases, and for individuals with health-care needs that are further down on the priority scale. Thomas Pogge, in fact, argues for *exactly this point*, in his extensive discussion of the incentives problem (Pogge 2008, Ch. 9). It's not enough to know that pharmaceutical R&D is very expensive, and that patents are better than other potential instruments for helping incentivize the massive resources it takes to get a new drug. Without an answer to these three questions (“Which drugs? Which diseases? For

whom?") we don't know enough to figure out whether pharmaceutical patents are functioning, defensible, and ultimately just.

When it comes to dissertations, graduate students' reach should exceed their grasp, or else what are postdocs and junior faculty positions for? The topics and issues covered here are only a small part of what a full treatment of IP would need to cover. The incentives problem is one of many outstanding issues raised, but not treated in depth, in this essay. What I've tried to do here is lay the foundations for a fuller, richer discussion of IP, one that incorporates the philosophical dimensions often neglected in IP theory and policy. An appreciation of these dimensions requires the conversation about IP, both amongst experts in law, economics, business, and politics, and all those individuals whose lives are touched by esoteric elements of IP law and policy, to be both broader and deeper. Broader, in incorporating the three axes discussed above. Deeper, in that the theory of IP must be embedded in discussions about what part government and law should have in regulating and shaping art, science, culture, and technology; about how to prioritize endeavors like biomedical research in a world with scarce resources; and ultimately about the role of institutions like IPRs in a just society.

Bibliography

Angell, M (1997) The Ethics of Clinical Research in the Third World. *New England Journal of Medicine* 337(12): 847-849.

Angell, M (2005) *The Truth About the Drug Companies: How They Deceive Us and What to Do About It*. New York, Random House.

Annas, GJ and Grodin, MA (1998) Human Rights and Maternal-Fetal HIV Transmission Prevention Trials in Africa. *American Journal of Public Health* 88(4): 560-563.

Attaran, A (2004) How do Patents and Economic Policies Affect Access to Essential Medicines in Developing Countries? *Health Affairs* 23(3): 155-166.

Attaran, A and Gillespie-White, L (2001) Do Patents For Anti-Retroviral Drugs Constrain Access to AIDS Treatment in Africa? *Journal of the American Medical Association* 285(15): 1886-1892.

Banerjee, A, Hollis, A and Pogge, T (2010) The Health Impact Fund: Incentives for Improving Access to Medicines. *The Lancet* 375(9709): 166-169.

Barlow, JP (1994) The Economy of Ideas. *Wired*, March.

Barron, A (2012) Kant, Copyright and Communicative Freedom. *Law and Philosophy* 31(1): 1-48.

Barton, JH (2004) TRIPS and the Global Pharmaceutical Market. *Health Affairs* 23(3): 146-154.

Becker, LC (1980) The Moral Basis of Property. In Pennock, JR and Chapman, JW (eds) *Property*. New York, NYU Press, 187-220.

Becker, LC (1981) *Property Rights: Philosophical Foundations*. London, Routledge.

Becker, LC (1992) Too Much Property. *Philosophy and Public Affairs* 21(2): 196-206.

Becker, LC (1993) Deserving to Own Intellectual Property. *Chicago-Kent Law Review* 68(2): 609-629.

Beitz, CR (1999) *Political Theory and International Relations*. Princeton, Princeton University Press.

Beitz, CR (2005) The Moral Rights of Creators of Literary and Artistic Works. *Journal of Political Philosophy* 13(3): 330-358.

- Benkler, Y (1999) Free as the Air to Common Use: First Amendment Constraints on Enclosure of the Public Domain. *New York University Law Review* 74(2): 354-446.
- Bessen, J and Meurer, MJ (2009) *Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators At Risk*. Princeton, Princeton University Press.
- Bhagwati, J (2002) Patents and the Poor: Including Intellectual Property Protection in WTO Rules Has Harmed the Developing World. *Financial Times*, September 17.
- Biron, L (2012) Public Reason, Communication and Intellectual Property. In Lever, A (ed) *New Frontiers in the Philosophy of Intellectual Property*. Cambridge, Cambridge University Press, 225-260.
- Blackstone, W (1765) *Commentaries on the Laws of England*. Available online through Yale Law School's library, at http://avalon.law.yale.edu/18th_century/blackstone_bk2ch1.asp.
- Boettke, PJ, Coyne, CJ and Leeson, PT. (2008) Institutional Stickiness and the New Development Economics. *American Journal of Economics and Sociology* 67(2): 331-358.
- Boldrin, M and Levine, DK (2008) *Against Intellectual Monopoly*. Cambridge, Cambridge University Press.
- Boyle, J (2010) *The Public Domain*. New Haven, Yale University Press.
- Boyle, J (2003) The Second Enclosure Movement and the Construction of the Public Domain. *Law and Contemporary Problems* 66(1/2): 33-74.
- Breakey, H (2009) Liberalism and Intellectual Property Rights. *Politics, Philosophy, and Economics* 8(3): 329-349.
- Buchanan, A, Cole, T and Keohane, RT (2011) Justice in the Diffusion of Innovation. *Journal of Political Philosophy* 19(3): 306-332.
- Calabresi, G and Melamed, AD (1972) Property Rules, Liability Rules, and Inalienability: One View of the Cathedral. *Harvard Law Review* 85(6): 1089-1128.
- Carlson, RH (2011) *Biology is Technology: The Promise, Peril, and New Business of Engineering Life*. Cambridge, Harvard University Press.
- Chakrabarty, AM (2002) Patenting of Life-Forms: From a Concept to Reality. In Magnus, D, Caplan, A and McGee, G (eds) *Who Owns Life?*. New York, Prometheus Books, 17-24.

Child, JW (1997) The Moral Foundations of Intangible Property. In Moore, AD (ed) *Intellectual Property: Moral, Legal, and International Dilemmas*. Lanham, Rowman and Littlefield, 57-80.

Coase, RH (1937) The Nature of the Firm. *Economica* 4(16): 386-405.

Collier, P (2008) *The Bottom Billion: Why the Poorest Countries are Failing and What Can Be Done About It*. New York, Oxford University Press.

Coyne, CJ (2013) *Doing Bad By Doing Good: Why Humanitarian Action Fails*. Palo Alto: Stanford University Press.

Crouch, RA and Arras, JD (1998) AZT Trials and Tribulations. *The Hastings Center Report* 28(6): 26-34.

Demsetz, H (1967) Toward a Theory of Property Rights. *American Economic Review* 57(2): 347-359.

Duffy, JF (2005) Intellectual Property and the Average Cost Isolationism Thesis. *Texas Law Review* 83(4): 1077-1096.

Dworkin, R (2002) *Sovereign Virtue*. Cambridge, Harvard University Press.

Easterly, W (2007) *The White Man's Burden: Why the West's Efforts to Aid the Rest Have Done So Much Ill and So Little Good*. New York, Penguin.

Fisher, W (2001) Theories of Intellectual Property. In Munzer, SR (ed) *New Essays in the Legal and Political Theory of Property*. Cambridge, Cambridge University Press, 168-200.

Flory, JH and Kitcher, P (2004) Global Health and the Scientific Research Agenda. *Philosophy and Public Affairs* 32(1): 36-65.

Flynn, S, Hollis, A and Palmedo, M (2009) An Economic Justification for Open Access to Essential Medicine Patents in Developing Countries. *Journal of Law, Medicine, and Ethics* 37(2): 184-208.

Gordon, WJ (1993) A Property Right in Self-Expression: Equality and Individualism in the Natural Law of Intellectual Property. *Yale Law Journal* 102(7): 1533-1609.

Griesemer, J (2011) Philosophy and Tinkering. *Biology and Philosophy* 26(2): 269-279.

Hegel, GWF (1991) *Elements of the Philosophy of Right*, Wood, AW (ed), Nisbet, HB (trans). Cambridge, Cambridge University Press.

- Helfer, LR and Austin, GW (2011) *Human Rights and Intellectual Property: Mapping the Global Interface*. Cambridge, Cambridge University Press.
- Heller, MA and Eisenberg, RS (1998) Can Patents Deter Innovation? The Anticommons in Biomedical Research. *Science* 280(5364): 698-701.
- Hess, C and Ostrom, E (2003) Ideas, Artifacts, and Facilities: Information as a Common-Pool Resource. *Law and Contemporary Problems* 66(1/2): 111-145.
- Hettinger, E (1989) Justifying Intellectual Property. *Philosophy and Public Affairs* 18(1): 31-52.
- Hobbes, T (1996) *Leviathan*, Tuck, R (ed). Cambridge, Cambridge University Press.
- Hughes, J (1988) The Philosophy of Intellectual Property. *Georgetown Law Journal* 77(2): 290-366.
- Hume, D (1987) Of the Original Contract and Of the Coalition of Parties. In Miller, EF (ed) *Essays Moral, Political, and Literary*. Indianapolis, Liberty Fund.
- Jaffe, AB and Lerner, J (2006) *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do About It*. Princeton, Princeton University Press.
- Johns, A (2011) *Piracy: The Intellectual Property Wars From Gutenberg to Gates*. Chicago, University of Chicago Press.
- Kant, I (1999) *Practical Philosophy*, Gregor, MJ (ed). Cambridge, Cambridge University Press, 1999.
- Kinsella, NS (2008) *Against Intellectual Property*. Auburn, Ludwig Von Mises Institute.
- Kremer, M (1998) Patent Buyouts: A Mechanism for Encouraging Innovation. *Quarterly Journal of Economics* 113(4): 1137-1167.
- Kremer, M (2002) Pharmaceuticals and the Developing World. *Journal of Economic Perspectives* 16(4): 67-90.
- Landes, WM and Posner, RA (2003) *The Economic Structure of Intellectual Property Law*. Cambridge, Harvard University Press.
- Lemley, MA (2005) Property, Intellectual Property, and Free Riding. *Texas Law Review* 83(4): 1031-1076.
- Lessig, L (2002) *The Future of Ideas: The Fate of the Commons in a Connected World*. New York, Vintage.

Lessig, L (2005) *Free Culture: The Nature and Future of Creativity*. New York, Penguin.

Litman, J (1990) The Public Domain. *Emory Law Journal* 39(4): 965-1024.

Locke, J (1988) *Two Treatises of Civil Government*, Laslett, PJ (ed). Cambridge, Cambridge University Press.

Locke, J (1997) Liberty of the Press. In Goldie, M (ed) *Political Essays*. Cambridge, Cambridge University Press, 329-338.

Lomasky LE (1987) *Persons, Rights, and the Moral Community*. New York, Oxford University Press.

Lomasky, LE (2007) Liberalism Beyond Borders. *Social Philosophy and Policy* 24(1): 206-233.

Lurie, P and Wolfe, SM (1997) Unethical Trials of Interventions to Reduce Perinatal Transmission of the Human Immunodeficiency Virus in Developing Countries. *New England Journal of Medicine* 337(12): 853-856.

Machlup, F (1958) An Economic Review of the Patent System. No. 15, US Government Printing Office.

Machlup, F and Penrose, E (1950) The Patent Controversy in the Nineteenth Century. *Journal of Economic History* 10(1): 1-29.

Mann, CC (2011) *1493: Uncovering the New World Columbus Created*. New York, Knopf.

Merges, RP (1999a) The Law and Economics of Employee Inventions. *Harvard Journal of Law and Technology* 13(1): 1-54.

Merges, RP (1999b) As Many as Six Impossible Patents Before Breakfast: Property Rights for Business Concepts and Patent System Reform. *Berkeley Technology Law Journal* 14(2): 577-615.

Merges, RP (2011) *Justifying Intellectual Property*. Cambridge: Harvard University Press.

Merrill, TW and Smith, HE (2007) The Morality of Property. *William and Mary Law Review* 48(5): 1849-1896.

Mill, JS (2008) *On Liberty and Other Essays*. New York, Oxford University Press.

Moore, AD (1997) Towards a Lockean Theory of Intellectual Property. In Moore, AD (ed) *Intellectual Property: Moral, Legal, and International Dilemmas*. Lanham, Rowman and Littlefield, 81-103.

- Mossoff, A (2007) Who Cares What Thomas Jefferson Thought About Patents? Reevaluating the Patent "Privilege" in Historical Context. *Cornell Law Review* 92(5): 953-1012.
- Mossoff, A (2012) Saving Locke From Marx: The Labor Theory of Value in Intellectual Property Theory. *Social Philosophy and Policy* 29(2): 283-317.
- Nagel, T (2005) The Problem of Global Justice. *Philosophy and Public Affairs* 33(2): 113-147.
- Nozick, R (1977) *Anarchy, State, and Utopia*. New York, Basic Books.
- O'Neill, O (1990) *Constructions of Reason: Explorations of Kant's Practical Philosophy*. Cambridge, Cambridge University Press.
- Ostrom, E (1990) *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, Cambridge University Press.
- Palmer, TG (1990) Are Patents and Copyrights Morally Justified?: The Philosophy of Property Rights in Ideal Objects. *Harvard Journal of Law and Public Policy* 13(3): 817-866.
- Plant, A (1934) The Economic Theory Concerning Patents for Inventions. *Economica* New Series 1(1): 30-51.
- Pogge, T (2007) Montreal Statement on the Human Right to Essential Medicines. *Cambridge Quarterly of Healthcare Ethics* 16(1): 97-108.
- Pogge, T (2008) *World Poverty and Human Rights*, 2nd ed. New York, Polity Press.
- Pogge, T (2009) The Health Impact Fund and Its Justification by Appeal to Human Rights. *Journal of Social Philosophy* 40(4): 542-569.
- Radin, MJ (1982) Property and Personhood. *Stanford Law Review* 34(5): 957-1015.
- Rand, A (1986) Patents and Copyrights. In *Capitalism: The Unknown Ideal*. New York, Signet.
- Raustiala, K and Sprigman, C (2012) *The Knockoff Economy: How Imitation Sparks Innovation*. New York, Oxford University Press.
- Rawls, J (1971) *A Theory of Justice*. Cambridge, Harvard University Press.
- Reeve, A (1986) *Property*. London, Macmillan.

- Reiss, J and Kitcher, P (2009) Biomedical Research, Neglected Diseases, and Well-Ordered Science. *Theoria* 24(3): 263-282.
- Rose, M (1995) *Authors and Owners: The Invention of Copyright*. Cambridge, Harvard University Press.
- Rosenberg, A (2004) On the Priority of Intellectual Property Rights, Especially in Biotechnology. *Politics, Philosophy, and Economics* 3(1): 77-95.
- Rosenberg, A (2012) Designing a Successor to the Patent as a Second Best Solution to the Problem of Optimum Provision of Good Ideas. In Lever, A (ed) *New Frontiers in the Philosophy of Intellectual Property*. Cambridge, Cambridge University Press, 88-109.
- Sachs, J (2006) *The End of Poverty*. New York: Penguin.
- Sachs, J and Malaney, P (2002) The Economic and Social Burden of Malaria. *Nature* 415(6872): 680-685.
- Satz, D (2010) *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. New York, Oxford University Press.
- Sell, SK (2001) TRIPS and the Access to Medicines Campaign. *Wisconsin International Law Journal* 20(3): 481-522.
- Schroeder, J (2006) Unnatural Rights: Hegel and Intellectual Property. *University of Miami Law Review* 60(4): 453-504.
- Shadlen, KC (2007) The Political Economy of AIDS Treatment: Intellectual Property and the Transformation of Generic Supply. *International Studies Quarterly* 51(3): 559-581.
- Shavell, S and van Ypersele, T (2001) Rewards versus Intellectual Property Rights. *Journal of Law and Economics* 44(2): 525-547.
- Shiffrin, SV (2001) Lockean Arguments for Private Intellectual Property. In Munzer, SR (ed) *New Essays in the Legal and Political Theory of Property*. Cambridge, Cambridge University Press, 138-167.
- Shiffrin, SV (2007) Intellectual Property. In Goodin, R, Petit, P, and Pogge, T (ed) *A Companion to Contemporary Political Philosophy*. Malden, Wiley-Blackwell, 653-668.
- Shiffrin, S (2008) The Incentives Argument for Intellectual Property Protection. In Grosserier, A, Marciano, A and Strowel, A (eds) *Intellectual Property and Theories of Justice*. London, Palgrave Macmillan, 94-105.
- Simmons, AJ (1992) *The Lockean Theory of Rights*. Princeton, Princeton University Press.

- Simmons, AJ (1998) Makers' Rights. *Journal of Ethics* 2(3): 197-218.
- Sreenivasan, G (1995). *The Limits of Lockean Rights in Property*. New York: Oxford University Press.
- Sreenivasan, G (2007) Health and Justice in Our Non-Ideal World. *Politics, Philosophy, and Economics* 6(2): 218-236.
- 't Hoen, E (2002) TRIPS, Pharmaceutical Patents, and Access to Essential Medicines: A Long Way From Seattle to Doha. *Chicago Journal of International Law* 3(1): 27-46.
- Tully, J (1983) *A Discourse on Property: Locke and His Adversaries*. New York, Oxford University Press.
- Vaidyanathan, S (2003) *Copyrights and Copywrongs: The Rise of Intellectual Property and How it Threatens Creativity*. New York, NYU Press.
- Waldron, J (1991) *The Right to Private Property*. New York, Oxford University Press.
- Waldron, J (1993) From Authors to Copiers: Individual Rights and Social Values in Intellectual Property. *Chicago-Kent Law Review* 68(2): 841-887.
- Wenar, L (2008) Property Rights and the Resource Curse. *Philosophy and Public Affairs* 36(1): 2-32.
- Williams, B (1973) A Critique of Utilitarianism. In Smart, JJC and Williams, B *Utilitarianism: For and Against*. Cambridge, Cambridge University Press.
- Wilson, J (2006) Microsoft on Copyright: An Ethical Analysis. *ICFAI Journal of Intellectual Property Rights* 5(4): 73-83.
- Wilson, J (2009) Could There Be a Right to Own Intellectual Property? *Law and Philosophy* 28(4): 393-427.
- Wilson, J (2012) On the Value of the Intellectual Commons. In Lever, A (ed) *New Frontiers in the Philosophy of Intellectual Property*. Cambridge, Cambridge University Press, 122-139.
- Wimsatt, WC (2007) *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, Harvard University Press.
- Yen, AC (1989) A First Amendment Perspective on the Idea/Expression Dichotomy and Copyright in a Work's 'Total Look and Feel'. *Emory Law Journal* 38(2): 393-436.
- Yen, AC (1990) Restoring the Natural Law: Copyright as Labor and Possession. *Ohio State Law Journal* 51(2): 517-559.