Chapter Three

Do We Own Our Data?

The Finders-Keepers Ethics of the Cyber Commons

James Goodrich

Amazon, Facebook, Google, and other companies collect information about us and our online activity. Using this information and data mining techniques, these companies are able to make good probabilistic inferences about what sorts of ads would be most effective in getting each of us to purchase particular goods. These companies are then able to sell ad space on their websites to other companies with the promise of maximizing the effectiveness of those other companies' ads. The collection of our personal data with the aforementioned business model has become a billion-dollar industry.

There is increasing concern, however, about whether this business model is morally acceptable. Some contest that it violates our rights to privacy, while others worry that it is unfair because it allows companies to internalize profits while externalizing their costs (Cheneval 2018). Still others believe that the model is exploitative or perhaps even constitutes theft (Arrieta-Ibarra et al. 2018). Whether using our data as a means of profit is exploitative or constitutes theft, however, surely turns on the question of whether we are entitled to the products of our data. In turn, this issue resolves itself into the question of whether we, in some sense, own our data in the first place (Ritter and Mayer 2018).

Much of our useful data is generated from our actions. In this sense, data looks to be the result of our labor. Thus, any account of property rights that closely links labor and our ownership entitlements will entail that we have rights to our personal digital data. This, in turn, would entail by the usual logic of property that so long as some consent mechanism is in place, companies would owe compensation to individuals for using their data. Failure to pay compensation would thus be constitutive of exploitation. (And if no consent mechanism was in place, this would constitute theft.) However, whether this rough line of thought is cogent hinges crucially upon the theory

of property rights we adopt and what we mean by capacious phrases like "the product of our labor."[1]

In particular, this chapter concerns itself with the specifics of the theory of property rights we might adopt. Which theory of property rights we adopt, after all, may well have implications about whether or not we have property rights in our personal digital data (Nolin 2019). We cannot, of course, do justice to all theories of property rights within a single chapter. We will need to focus more narrowly. My interest is in which implications the Kirznerian theory of property rights has for the debate over property rights in personal digital data. The results of this inquiry, however, should nevertheless prove instructive for teasing out the implications of other theories of property rights in the context of personal digital data.

To foreshadow my conclusions, the Kirznerian theory of property rights *itself* does not have the tools to claim that each of us as individuals have rights to our personal digital data (or rights to the profits generated by the use of that data). However, the larger Kirznerian view of distributive justice may give us reasons to prefer a set of institutions that would recognize conventional versions of such property rights. Roughly, the thought is this: data should be considered a common pool resource. Rights to access a commons (in order to acquire property rights over some resource in that commons) differ from rights to acquire private property in the first place (Schlager and Ostrom 1992). The Kirznerian's wider views of distributive justice may give us reasons to adopt institutions that constrain access to the cyber commons of personal digital data. Respecting such constraints imposes at least *some* duties on Amazon et al. that are owed directly to the individuals who could make use of that commons. These duties, in turn, imply some correlative rights for these individuals to their digital data (or profits from that data).

Here is how I will proceed. First, I will consider the question of whether we can own our data really makes good conceptual sense. Some have expressed skepticism on this point. Next, I will briefly introduce the Kirznerian theory of property rights and raise some worries about one attempt to bring the Kirznerian theory into conversation with disputes over our data. This will lead to an examination of each part of the Kirznerian theory, which I will argue provides no help to justifying claims about natural property rights in our digital data. Finally, I will argue that Kirznerians have available to themselves a conventional account of property rights that would have interesting implications for the debates surrounding ownership and our data.

CONCEPTUAL PRELIMINARIES

I have assumed that it made sense to ask, "Do we have property rights over our personal digital data?" But does it? Some have contested that, due to some apparent conceptual or metaphysical difficulties, personal digital data is just not the kind of thing that could, even in principle, be owned. Our inquiry is thus a nonstarter if these worries are on mark. As such, it behooves us to discuss them here at the outset before we turn to the questions of moral theories that presuppose their answers. (However, if you find yourself antecedently convinced that data could in principle be owned, then you may wish to skip this section.)

Let's start simply. Some kinds of things clearly *can* be owned. I own the laptop I'm typing on, the chair I'm sitting on, and the coffee mug out of which I'm drinking. Other kinds of things clearly cannot be owned. I could not, even in principle, own the number nine, the color green, or the second law of thermodynamics. But what distinguishes the kinds of things that can be owned from the things that cannot? And on which side of this line does our personal digital data fall? The first set of objects I named—the laptop, the chair, and the coffee mug—are concrete particulars. They are among the physical stuff we can bump into in the world. The second set of objects I named—the number nine, the color green, and the second law of thermodynamics—are abstract universals. We cannot bump into these things. And more importantly, we cannot do a lot of other things with abstract universals. We cannot sell them, earn profits from them, or exclude others from using them. All of these characteristics are hallmarks of the kinds of things we are capable of owning. Thus, the fact that abstract universals seem to lack these features gives us reason to doubt that they could be the proper objects of ownership.[2] One might look at our personal digital data and think, "Well, I can't bump into that either." So perhaps our data is more like the aforementioned abstract universals than it is like the laptop, chair, or coffee mug. And if so, perhaps our personal digital data is also not the kind of thing that can be owned, even in principle.

Of course, we can in principle sell our data, earn profits from it, and exclude it from being used by others. Suppose there was no monitoring of our online activity. But instead, I simply wrote down on a piece of paper some demographic information about myself, all the Google searches I made in a day, and all of the links that I clicked on. I could, at least in principle, sell this piece of paper or prevent you from looking at it. This may not mean that I have the right to do so. But that is a moral point, not a conceptual or metaphysical one. In other words, all our personal digital data is just a kind of information, and we're all familiar with ways to sell and exclude people from

information. Entire businesses are founded on the very idea that we can exclude information from those who do not pay for it—bookstores and universities are clear examples. But this might extend to any sector of the economy in which one pays for the advice of experts on some question or another since such a service involves an exchange of information to the buyer. Insofar as our data is information, it seems reasonable that despite any conceptual or metaphysical worries regarding our data as property, it is not that such data is, in principle, incapable of being sold or that we cannot, in principle, exclude others from using it.

Here is a different worry.[3] The worry is that our personal digital data cannot be the object of a property right because information about one individual is often also information about another individual. For example, my mother's birthday is data about my mother at least as much as it is data about myself. It is *her* birthday, but it is also *my* mother's birthday. And it's also my sister-in-law's mother-in-law's birthday, her mother's daughter's mother-in-law's birthday, and so on. Who has a right to this data? What principle could possibly tell us?

These questions may appear intractable. However, they are not. As a preliminary conceptual matter, "Ellen Goodrich's birthday" and "my mother's birthday" refer to the same date. But they are nevertheless not identical pieces of information. After all, some facts could be true of a person whose mother's birthday was *X* that are not true of the mother who is having the birthday. Here is a relevant example. My consumption habits approaching my mother's birthday are likely to be quite different than my mother's consumption habits. I will want to buy a present for my mother, and my mother might hold off on buying something because she believes that I might give it to her for her birthday. Thus, all we need to see is that the referent—namely, the birth date in question—does not exhaust the content of the data or the uses to which it could be put.[4] Therefore, we should opt for a relatively fine-grained notion of "personal digital data" and a fine-grained understanding of our rights over that data. If I own my data, I thus own the data that "my mother's birthday is *X*." My mother owns the data that "her birthday is *X*." There is at least no conceptual conflict here.

But perhaps a more practical problem remains. The problem is this: suppose I sign a contract with Facebook to give them all of my personal digital data. Among the data they will receive is "my mother's birthday is *X*" and "my mother is Ellen Goodrich." From this data, they can infer that "Ellen Goodrich's birthday is *X*." Thus, it's difficult for me to sign away my data in a contract to Facebook while preserving my mother's ability to successfully exclude Facebook from using her data. Put somewhat differently, they may not need to use my mother's data if they can use mine to target ads toward

her anyway. Thus, the objection goes, the fine-grained information response I provided above does not avoid the practical problem of what to do when two individuals' data is connected.

However, this practical problem is not the hurdle it may first appear to be. What it means is that one of the entitlements associated with my property right to my data—my ability to consent to others using my data—is limited by the rights of other individuals. This is hardly surprising. In general, I have a right to bodily autonomy. I can move and use my body as I like. Perhaps I like to shadow box. However, if someone stands in front of me in a public place where I usually shadow box, I may not intentionally shadow box there at that moment. Why? Because what would usually be an exercise of my right to control the movements of my own body now involves violating the rights of others, and I clearly am not allowed to do so. Thus, I'm simply not allowed to sign a contract that would give some of the data I usually have rights over to Facebook if this would violate my mother's rights. In practice, this might mean that I could contract out either the data that "my mother's birthday is *X*" or "my mother is Ellen Goodrich" but not both. Perhaps in practice this would also mean that someone would need to create an app that could check for logical or probabilistic entailments of one's data to ensure no such violation occurs.

An alternative solution would be to only recognize moral rights to certain predefined bundles of data so that no such conflict would occur. This might mean, in practice, that some natural moral rights to data may not be recognized, but this may be an acceptable consequence if it provides a better overall solution to the practical problem. What would these predefined bundles look like? I couldn't tell you. They'd have to be determined via a political process. The bundles would thus, in some sense, be socially constructed. This need not bother us, however, since we have property rights over socially constructed objects already (e.g., currency, shares of funds, and so on). I thus conclude that my guiding question—Do we own our personal digital data?—makes conceptual sense. I can make sense of it, and I suspect you can too.

KIRZNER'S THEORY OF PROPERTY RIGHTS

Most of us believe that what is just or fair depends upon, *inter alia*, individual moral rights. A subset of these moral rights are ownership rights—a bundle of entitlements to (1) use an owned good, (2) exclude use of the owned good by others, (3) rent or sell the use of the owned good to others, and (4) gain profits from the use of the owned good. Thus, if you own a car, *ceteris paribus*, I owe you a *pro tanto* moral duty to not use your car without your permission.[5]

If I did use your car without your permission, then I would, *ceteris paribus*, be violating your rights and thereby wronging you. And wronging you, in turn, would be unfair or unjust. (Though to be clear, the dimension of justice associated with property rights may not exhaust all dimensions of justice nor morality as a whole.)

Most theories of the content, scope, and grounds of ownership rights have been developed with tangible goods in mind. Given the historical context in which such theories were developed, this made perfectly good sense. In recent decades, however, questions concerning ownership over non-tangible goods have become increasingly important politically, socially, and economically. Perhaps the most well-attended to example of putative ownership rights in non-tangible goods is intellectual property. Of course, what we should say about intellectual property, especially from the perspective of moral theory, is far from settled. But questions concerning intellectually property only make up a small subset of potential ownership rights in non-tangible goods. There are less well-attended to moral and political questions concerning potential ownership rights in other kinds of non-tangible goods. This chapter focuses on one such family questions: What are the contents, scope, and grounds of ownership rights in digital data—if there are any at all?

As I said in my introductory remarks, our focus will be on what the Kirznerian theory of property rights should say about digital data. To a first approximation, the Kirznerian theory of property rights is a version of the Lockean theory of property rights.[6] We can state the Lockean theory, roughly, as follows:

> **The Lockean View**: An agent S possesses an ownership right over a good G if and only if (i) G is a part of S's person OR (ii) G is the result of S's legitimate appropriation AND (iii) enough and as good of G is left for others.

There are two ways appropriation of a good might be legitimate. Either S receives a good G from someone who owns G or S does not. If someone owns G, then G must be given to S by contract or consent of its owner. However, if G is unowned, then S owns G if S meets *an original appropriation condition*. What exactly constitutes the original appropriation condition in Locke's original view is debated, for it turns on what exactly Locke meant by "mixing one's labor." As Robert Nozick famously made the point, one does not own the ocean merely by mixing orange juice one has made with it (Nozick 1974).

Contemporary neo-Lockeans agree with much of Locke's picture, but they often disagree about how we should best understand the original appropriation condition. (They also often disagree with how to understand clause (iii), which is sometimes called "The Lockean Proviso." More on this in due course.) This is where Kirzner comes in. Israel Kirzner offers a novel *suffi-*

cient, but not necessary condition for what it takes to originally appropriate a good. His thought is as follows (Kirzner 2016). Most of us would agree that if one causes there to exist some new good G in excess of the good G* required to create G, then, *ceteris paribus*, one is entitled to at least the difference in value of G and G*. Let's call this "the Creation-Entitlement Premise." The Creation-Entitlement Premise is itself a sufficient condition for legitimate original appropriation of some good. Kirzner exploits the Creation-Entitlement Premise by offering a novel interpretation of the "causes there to exist some new good" locution. His key insight is that *discovery* of a new good is itself a form of causing a new good to exist, *practically speaking*. That is, there are epistemic constraints on what we should consider the total set of *status quo* owned or unowned goods to be. Put roughly, if no one is aware of a good, then no one can make use of it. If no one can make use of it, it does not exist, practically speaking, and therefore no one is entitled to it.[7] Thus, when someone discovers a valuable item, they are in fact causing the total set of owned or unowned goods to increase. They are "enlarging the pie," as it were. Discovery, thus, is an act of creation according to what we might call Kirzner's "Finders-Keepers" version of the Creation-Entitlement Premise, or "the Finders-Keepers Premise" for short. Our question for now becomes: What, if anything, does the Finders-Keepers Premise imply about our personal digital data?

AN INSTRUCTIVE MISUNDERSTANDING

It will be instructive for our purposes to consider a misunderstanding of the Finders-Keepers Premise already present in the literature. Marijn Sax has argued that (1) something akin to Kirzner's Finders-Keepers Premise indeed underlies much of the rationale for the potential moral entitlement of companies like Amazon, Facebook, and Google to use our data and (2) Kirzner's Finders-Keepers Premise can be shown to possess problematic assumptions in the case of our personal digital data (Sax 2016).

With respect to (1), Sax argues that the aforementioned companies are best understood as discovering a valuable product of our data because our data is not itself so valuable until one exploits data-mining techniques to extract further information from it. And in this sense, he believes Kirzner's Finders-Keepers Premise applies. This is an interesting point. Notice that any version of the Creation-Entitlement Premise that was non-epistemic—to wit, that required the capital "E" Existence of a new valuable item—likely would not apply to the case of Amazon et al. Why? Well, new information is plausibly not created by the act of data mining since that information is entailed by the

underlying data. Or it would be, at any rate, highly controversial to say that new, valuable objects were *created* as opposed to *discovered*. The resulting information collected from data mining our personal digital data is, in other words, already out there to be discovered. And on Kirzner's discovery interpretation of creation, this is good enough to create an entitlement.

Now, I would like to note a minor, passing objection to (1). No one to my knowledge actually employs Kirzner's thinking in the way described by Sax, so Sax risks strawmanning both the defenders of current data practices and Kirzner by presenting them as allies.

Let's now consider (2). Sax deploys a combination of objections to the Kirznerian Finders-Keepers Premise approach to current data practices, but the heart of his objection seems to rest on the idea that the Finders-Keepers Premise, if it's applied, leads us astray in the case of personal digital data. Why? Here he appeals to the idea that our personal digital data is somehow closely connected to our person: "the idea that your identity as a person is always necessarily constituted—at least partly—by your information" (Sax 2016, 29). Sax seems to be implicitly invoking clause (i) of the Lockean View. Just as we own our bodies, so too do we own our information. Thus, according to Sax, just as a company cannot use our bodies without consent to derive entitlements to profits, so too they cannot use our information.

Sax makes a rather bold metaphysical claim in saying that our information is necessarily constitutive of us. To see the falsity of Sax's claim, it's sufficient to see that not all of our information is necessarily constitutive of us. Indeed, likely very little of it is. Whether I used Google to search the terms "Sax," "Kirzner," and "big data" is hardly necessarily constitutive of me. I would have been the exact same person typing this chapter had I never performed that action. Moreover, whether I was born in Generation Y or Generation Z may be relevant data. However, such a difference could hardly be constitutive of me. After all, which of the two generations I was born into comes down to whether I was born a minute before or a minute after midnight on December 31, 1994. But a mere difference of two minutes could hardly make me a different person. It would likely indicate more about the procedures of the given hospital at which I was born than it would about me.

There are, of course, moves Sax could make in response. But the further we wade into such metaphysical waters, the more controversial and therefore less compelling Sax's objection is liable to become. If Sax must rely on a much more complicated story about how quite a lot of our data that does not seem essential to us in fact is essential, we should begin to wonder whether we are adopting a less plausible metaphysical picture to gerrymander the result we want about data. And surely, that is not the direction in which we should go.

For the sake of argument, however, suppose Sax is right that our personal digital data is necessarily constitutive of us. Sax's objection nevertheless fails to be compelling. Kirzner's Finders-Keepers Premise is just a sufficient condition for legitimate original appropriation of some good. Thus, it tells us nothing about what should be said in cases that involve the use of already-owned property. In other words, Kirzner's view is neo-Lockean. Therefore, it matters to him a great deal whether a given good is already owned. Kirzner would not think, for example, that if Jack stole a bunch of apples from Jill's orchard and made apple juice out of them, then Jack owns the resulting apple juice even if no one had previously discovered the concept of apple juice. Therefore, whether Kirzner's Finders-Keepers premise is even implicated in the case at hand depends on the prior question of whether or not individuals own their data. And thus, Sax's point about (1), while interesting in its own right, is only relevant if Kirzner's view is antecedently committed to the view that individuals do not own their own personal digital data.

Sax might respond by saying that he was never attacking the Kirznerian view, and he was just attacking its application to personal digital data. He could certainly be clearer on this point. Again, framing the issue as he has creates at least the specter of a strawman. In any case, the Kirznerian should be interested in the question, "Do individuals own their personal digital data?" If they do, then Amazon et al. are wronging individuals. If they do not, then perhaps Amazon et al. are not simply failing to wrong individuals; perhaps they are entitled by the Kirznerian's lights to the profits of the products of everyone's personal digital data.

THE LOCKEAN PROVISO

The Kirznerian is working within a neo-Lockean framework. One way to understand Sax—once we have clarified that his objection is not actually to the Kirznerian *per se*—is that Sax was suggesting that we do own our personal digital data because of clause (i) of the Lockean View. For reasons that I previously discussed, I rather doubt that this is a successful move. A different move one might make is to argue on the basis clause (iii) of the Lockean View—also known as "the Lockean Proviso"—that individuals own their own personal digital data. Notice, however, that the Lockean Proviso is a constraint on the Creation-Entitlement Premise. Thus, it could only be invoked to rule out the acquisition of some kinds of property. (In other words, it's a necessary condition of a sufficient condition of property acquisition.) It cannot be invoked to actually ground a property right itself. It is thus of no use

if what we want is a theoretical story about how people possess a legitimate property right; it can only tell us who does not have such a right. Because it could become relevant, however, it's worth pausing to consider whether the Lockean Proviso might rule out the putative ownership right by Amazon et al. over our personal digital data.

Kirzner himself makes some interesting remarks on how we should best understand his Finders-Keepers Premise in relation to the Lockean Proviso. For Kirzner, insofar as the Finders-Keepers version of the Creation-Entitlement Premise is in play, the Lockean Proviso may not apply (Kirzner 2016, 149–55). Why? In Kirzner's interpretation, the argument for the Lockean Proviso assumes that there is a fixed amount of resources. When individuals make a useful good, they have to use up other resources to make it. Given that resources are scarce, this means that for a system of property rights not to lead to horrible consequences for many, it has to be that individuals have an equality opportunity to make use of the scarce resources. If the system lacks this, the argument goes, then those individuals without an equal opportunity to make good use of a scarce resource will be harmed by being deprived of various resources. Individuals have a right to their bodies and against being harmed. Kirzner, however, argues that the kind of entrepreneurial discovery at work in the Finders-Keepers Premise invalidates this argument. The reason is that entrepreneurial discoveries are, in a certain sense, creations *ex nihilo*. They are discoveries of more value than we previously knew were there. Thus, discovery in Kirzner's view increases the size of the value pie. And if that is right, then no one can be said to be harmed by the discoverer's acquisition of that value (or perhaps means to it). This is because it does not impose the potential harms of deprivation on anyone for the discoverer to acquire that extra value. Thus, the thought goes, the Lockean Proviso does not apply if the Finders-Keepers Premise does.[8]

Whatever one thinks of this argument on Kirzner's behalf, it must be admitted that the Lockean Proviso applied to Amazon et al. would conflict with the heart of the Kirznerian position on discovery. And if that is so, then any invocation of the Finders-Keepers Premise could not also invoke the Lockean Proviso in defense of ownership rights (in anything). In other words, a Kirznerian can only invoke the Lockean Proviso by not invoking what is distinctive of the Kirznerian position. And if they cannot invoke what is distinctive of the Kirznerian position, then there is no way the Kirznerian position could rule out the possibility that Amazon et al. have a given property right. Indeed, as I mentioned in the previous section, the Finders-Keepers Premise can be used—if we agree with Sax's point (1) —to defend Amazon et al.'s property right in our data (or to the use of our data to make a profit).

THE FINDERS-KEEPERS PREMISE

At this point, I have argued that neither clause (i) nor (iii) aids us in explaining why or how individuals could have property rights over their personal digital data. This leaves clause (ii), which, interestingly, is what the Kirznerian version of the Lockean View is about. Put more precisely, it is about original appropriation. Because consensual appropriation of good G presupposes ownership of G, only original appropriation is relevant to the issue of whether each of us owns our personal digital data. This invites an interesting question: Could the Finders-Keepers Premise itself underlie our initial ownership rights to personal digital data, thereby undermining the unrestricted rights of Amazon et al. to profit from our data?

Notice, of course, that other versions of the Creation-Entitlement Premise might be able to play the aforementioned right-grounding role as well. And it is certainly open to the Kirznerian to invoke such a premise. However, all such versions of the premise are controversial, so considering such positions would take us far afield of what the Kirznerian is entitled to by their own lights.

One candidate proposal is that our data-generating actions constitute discoveries of our own preferences. What we are disposed to prefer may in one sense be "out there," but it's not until we perform data-generating actions that we discover what our own preferences are. I have three worries about this view.

My first worry is that some of my personal digital data would not be covered by this view, even in principle. For example, it may well be useful for a business to know when my birthday is without knowing anything else about me. And yet, just having a birthday is in no way a discovery. Or take another example. Suppose I become friends on Facebook with someone I have known for a very long time outside of the internet. It is difficult to see how I have learned anything about my preferences by agreeing to be their Facebook friend. And yet, this is potentially the kind of data at stake in our larger discussions about property rights and personal digital data. Of course, this might mean we have to simply restrict the scope of the Kirznerian response.

The second worry about this proposal is that in performing many of our data-generating activities, we do not recognize the results of those activities as teaching us about our preferences. This is a slightly different worry from the first. When I used Google to look for Sax's article, I did not learn anything about my preferences. Or perhaps I did. But what I learned surely is not apparent to me. And yet, this was a data-generating action. Again, it's open for the Kirznerian, of course, to claim that we only have a right to the data resulting from activities that do teach us about our preferences. This may be right

as far as morality goes. However, it is worth noting that it would be a difficult distinction to codify into law, thus creating potential problems downstream from this discussion for how to handle Amazon et al. in practice.

The third worry is that this view would overgenerate data ownership if the concept of a "data-generating" activity is sufficiently wide. Do I have ownership over the contents of my perceptual experiences or the visual contents of a photograph if I find that I enjoy looking at sunsets? After all, such contents are plausibly just information. Perhaps we do have such ownership rights, but this is surely taking us in a direction one would not have thought we would end up going in when we started with the Finders-Keepers Premise.

However, even if we grant that these three worries can be abated, it is unclear whether a discovery of my own preferences counts as the kind of discovery that is at the heart of Kirzner's theory. Kirzner was interested in entrepreneurial discoveries. These are discoveries of new valuable uses for token resources, services, or assets. Thus, even if I learn something about my preferences via my data-generating activities, it is far from clear that what I am discovering is a new, valuable use for my data. For example, suppose I learn that I enjoy chickpea vindaloo. This might be valuable information in one sense. There are now more entrée options I enjoy than I previously thought. But it is unclear that this particular discovery generates new value by itself. If I learn I can make chickpea vindaloo at a lower cost than any other dish that I equally enjoy, then this might generate some new value (for me at least). If I instead discover that chickpea vindaloo is always more expensive than dishes I equally enjoy, then this is not a value-increasing discovery. Thus, there are at least some differences between discovering I have a preference and discovering a piece of information with value. They can, of course, be related in important ways. But such information is insufficient without the right mean-ends relationship in place.

The Kirznerian does not seem to have resources to claim, by the light of their neo-Lockean commitments at least, that we have ownership rights in our digital data. Despite this seeming skepticism, I will now shift gears to argue that there are limits to the ownership rights of companies like Amazon, Facebook, and Google in our digital data. Properly understood, these limits imply duties on the part of Amazon et al. to the rest of us. We can claim that such duties do correlate to rights possessed by individuals. However, these rights are not exactly to their personal digital data per se. And all of this rests, with some exegetical liberties taken, on a wider understanding of Kirzner's Finders-Keepers ethics.

NATURAL VERSUS CONVENTIONAL RIGHTS

Notice that Lockean commitments, as they are typically understood, are only commitments to so-called "natural" or "nonconventional" rights. This leaves it open for the Kirznerian to argue that recognizing a conventional property right to personal digital property could be morally justified on consequentialist grounds. This involves two steps. The first step is identifying a good consequence that a society could legitimately promote. This requires attending to the usual squabbles of mainstream liberal political philosophy. But it could be done. The second step would be to show empirically that enforcing a conventional right to personal digital data would promote the aforementioned good. Again, this could be done. And if it could, the Kirznerian could defend the position that personal digital data should be recognized as owned by those who produced the data even if this is not a matter of natural justice. Moreover, this would be in keeping with a Kirznerian view insofar as the thing to be promoted looks valuable by Kirzner's lights. For example, if one of two schemes of property rights looked to better promote entrepreneurial discovery, one could see how Kirzner's wider conception of Finders-Keepers ethics might entail that the property rights regime that better promoted entrepreneurial discovery is preferable.

But there is really a third step as well. If we think that the heart of Sax's point (1) is correct, then by the Kirznerian's lights, Amazon et al. satisfy the Finders-Keepers Premise. Insofar as that is true, Amazon et al. have an entitlement to the products of our data assuming we do not naturally own them. Furthermore, the good consequences of recognizing conventional ownership rights in personal data would need to be sufficiently weighty to outweigh Amazon et al.'s natural entitlement to the profits our personal digital data. Thus, the Kirznerian must find a sufficiently weighty good to be achieved, lest they allow for natural ownership rights to be frequently outweighed by lesser, conventional consequentialist considerations. And this may not be in keeping with Kirzner's distinctive claim that pure profits in capitalism are, in fact, just.[9]

However, there may be a different strategy. This would be to show that Sax's point (1) is wrong or that insofar as it's right, it's limited in scope. In other words, Amazon et al. may have some moral entitlement to profiting from our data, but they are currently profiting more than they are entitled to by the Kirznerian's lights. If that is so, one could explain why some have the intuition that there is something morally problematic about Amazon et al.'s business model while still making essential appeals to some claims about the Kirznerian view of property rights. In the next two sections, I will sketch one promising way this could be accomplished.

DISCOVERY IN THE CYBER COMMONS

The Kirznerian should distinguish between the discovery of valuable uses of some asset and the discovery of a commons producing those assets. It is plausibly the former and not the latter that generates a property right. To see why, think of our data-generating activity as a kind of commons. Our continual acts of searching, posting pictures, and so on should thus be thought of similarly to more familiar self-sustaining natural resource commons such as fisheries. Just as a local environment will create a self-sustaining system in which new fish will continually be produced, cyber commons also continually produce new data, which can then be put to various valuable purposes. This intuition can be extended by recognizing a layer of separation between us and what is useful about our data: it is useful in aggregate to spot trends. It is the emerging patterns to which we contribute that allow for serious companies like Amazon et al. to flourish.

Let's note a few further things about the cyber commons and its relationship to the Kirznerian's more general Finders-Keepers ethics. First, the Kirznerian view of property rights can easily explain why the discovery of new valuable uses of data can entitle the discoverer to profits from those *token* uses. But Kirzner would never hold that the discovery of the uses of data could entitle an individual to the profits generated by *types* of valuable uses. For example, if I discover that you can make juice out of apples, this entitles me to profit from the juice I make. However, *ceteris paribus*, I am not thereby entitled to profits from the juice that others might make from the apples in their orchards.[10] It's thus the discovery of token uses of resources that deliver entitlements, not types of uses. If this is right, it in turn explains why many different companies can act justly (or at least not unjustly) when they use our personal digital data. No one company can have the entitlement to the big data business model itself.

Second, the question of who is entitled to access the cyber commons in the first place differs from the question of whether those who discover valuable uses of data in the cyber commons are entitled to the profits of their discovery. Again, let's return to a natural resource commons like a fishery. Individuals who catch fish in these fisheries may well be entitled to the profits generated by selling those fish. However, there are often mechanisms in place to ensure that no one individual gets complete, unfettered access to the commons (Ostrom 1990). Like Kirzner's view, nothing in the logic of a neo-Lockean view of property rights entails that an individual has a right to (or a duty against) accessing a commons. This creates a kind of elbow room for the more general Kirznerian Finders-Keepers ethics. While nothing in the neo-Lockean Finders-Keepers Premise entails that each of us presently pos-

sesses a property right in our data that Amazon et al. is violating, it is unclear whether Amazon et al. have an initial right to access the cyber commons constituted by our data. Perhaps it is in keeping with Kirznerian Finders-Keepers ethics more generally that such access rights to the just acquisition of property can or should be withheld from Amazon et al. We will probe this possibility further in the next section.

Third, the cyber commons clearly differs from a natural resource commons in a variety of ways. Perhaps the most fundamental difference is that the cyber commons is thoroughly socially constructed. This is true in the simple sense that, without human beings engaging in social interactions, we would not have this cyber commons.[11] But the cyber commons is also socially constructed in the further sense that data-generating activities are both constituted by and constitute cultural activities. In other words, many of our data-generating activities are only economically valuable insofar as they are enmeshed within various cultural institutions. I am engaging in social activity when I search for concert tickets or "like" a post about my friend's new job. I am performing actions that reflect my tastes in cultural activities (e.g., I have bought concert tickets and not theater tickets) and signal some facts about who is within my social circles. Moreover, insofar as market activities are cultural and social activities, there is a credible claim to be made that Amazon purchases are themselves a kind of social and cultural activity.[12] The data generated from these activities is thus about me, my social relations with others, and my preferences. And it is the continual dynamic updates to this kind of information—often within a larger data set, of course—that have economic value. Natural resource commons, like fisheries, are occasionally sites for social exchange, but they do not possess these features by necessity and perhaps rarely possess them saliently.

Moreover, such data-generating activities spawn new cultural artifacts and thereby support new social constructions. That is, "internet culture" is itself a kind of social product of our data-generating activities. People create and share "memes" more efficiently than they once did. And these memes then get turned into products. One only needs to type "grumpy cat" into Amazon.com to find many pages of products inspired by an internet joke. Buying and wearing a Grumpy Cat T-shirt would do more social signaling than my selling of a fish to a cafe would. My point here is simply that there is a kind of social feedback loop to our data-generating activities. We are simultaneously engaging in and creating widespread cultural artifacts by engaging as we do in online activities.

This point about the social construction of the cyber commons may seem like little more than a digression. However, it is emphatically not. One obstacle to appreciating the idea that our data-generating activities could constitute

a kind of commons is that we are used to thinking about particular commons in terms of asocial natural resources. But there is no conceptual reason we must think this way. And if the thought that our data-generating activities are a kind of commons is correct, then we must ask the question of who has rights of access to this commons. Our answer to this question is important if we are to consider the wider question: What, if anything, does the Kirznerian have to say about whether we should recognize individual data-generators as having property rights in their personal digital data?

We can make this point about the relevance of the social construction of the cyber commons to the Kirznerian's Finders Keepers ethics clearest by considering a worry about the direction in which I am heading. One might say,

> These are all fine points. However, I cannot see how they could become relevant. Clearly Amazon et al. have rights of access to the cyber commons. After all, this cyber commons was *created* for us by various companies. It's not as if they simply came upon it like a new body of water in which they began fishing. Amazon, Facebook, and Google spent a great deal of time and capital to construct their website and make them user-friendly. So how could any of this make a difference? It's clear how the aforementioned rights to access fall.

Not so fast. It's unclear whether the assumption that Amazon et al. created the cyber commons is in fact true. Here is a different interpretation: Amazon et al. improved the cyber commons. After all, our data and our data-generating activities are nothing new in one important sense. Our data—or our most valuable data—is about and generated by our social or cultural exchanges. But we've always had such exchanges. And there has always been information about such exchanges. What Amazon et al. has allowed us to do (among many other things, of course) is engage in cultural and social exchanges at lower costs to ourselves. To be clear, the improvements made by Amazon et al. may well be moral entitlement-generating improvements. That is, we may owe Amazon et al. something for reducing the costs of some of our cultural and social activities. However, this does not imply that they have the unfettered and unrestricted property rights in the cyber commons that would be generated if they had created something new as opposed to improving something that was already there.

Consider an analogy. Suppose there is an unowned fishery that a given community makes use of. Someone comes along and improves this fishery, and it thus produces more fish. *Ceteris paribus*, the community that uses this fishery benefits from this improvement. All the fishers can acquire more fish to sell. Intuitively, the person who improved the fishery is morally entitled to some nontrivial compensation for their actions. Perhaps they're even owed

some percentage of the improved output of the fishery. Suppose, however, that they are unwilling or unable to do some of the labor of fishing itself. If that is right, there is some percentage of the improved output of the fishery that would not (intuitively) be owed to this person. Moreover, this person would not be owed a percentage of the output of the fish that were already being produced by this fishery.

In other words, as I mentioned when outlining the basics of a neo-Lockean theory of property rights, most of us would agree that if one causes there to exist some new good G in excess of the good G* required to create G, then, *ceteris paribus*, one is entitled to at least the difference in value of G and G*. There are therefore intuitive limits to what one is entitled to by virtue of improving a commons. If this analogy is apt, we might say the same of Amazon et al. and the cyber commons. There are limits to that which they are entitled. And perhaps they are not entitled to all of their profits.[13]

Let me now put my larger point more succinctly. The discoveries of valuable uses for data made by companies like Amazon et al. entitle them to at least some of the resulting profits. However, whether this entitles them to the present distribution of profits acquired from the collection of our data is a separate matter. Insofar as their collection of our data based off of our activity online is a way for them to capitalize on what was already there—our larger cultural and social exchanges—the cyber commons should be seen as an improvement on the already-existing cultural and social commons. If that is right, then they do not morally possess the entitlements to all of the profits they now have.

I have argued that if it's best to construe Amazon et al. as having improved the cyber commons and not created it, they *might* be entitled to less than the status quo legal situation would suggest. Thus, Sax's point (1) is, strictly speaking, true, but it is more limited in scope than we might have previously thought. However, just because Amazon et al. lack this entitlement does not mean that they have done anything morally wrong. Some moral reason for them to not act as they have would need to be shown. Or alternatively, there would need to be a moral reason to institute a different set of restrictions on Amazon et al.'s use of the cyber commons. In the next section, I will take up this question.

GOVERNING THE CYBER COMMONS

I argued in the previous section that in some interpretations of what Amazon et al. did for the cyber commons, they are not entitled to all of the profits of using our data. This does not mean they have done anything wrong. And for

the sake of argument, I will set aside the question of whether they have. What I am interested in now is that, given that Amazon et al. lacks an entitlement to at least some of their profits, are we entitled to limit their access to the rest of the remaining profits? This will depend upon what else is at stake. If nothing else of moral relevance is at stake, then it would constitute little more than spite to deny Amazon et al. the additional profits. As I said before, the Kirznerian could claim here that, so long as it does not run afoul of other morally important principles, we could seek to promote entrepreneurial discovery. In the Kirznerian picture of the world, this is important; for it is the entrepreneur at rock bottom who explains the success of the market process. It's the entrepreneur who grows the value pie. Our question thus becomes this: Can we give to Amazon et al. what they are owed for their improvements to the cyber commons while also better promoting entrepreneurship within the cyber commons?

Any set of institutions that attempts to do so would be wise, other things being equal, to avoid top-down control by a bureaucracy. Such control would allow bureaucratic forces to shape who has access to the cyber commons. This, in turn, may lead to corruption due to political capture, a deficit of entrepreneurship, or an objectionable distribution of access to entrepreneurial opportunities. And, of course, which form the best bottom-up scheme should take is largely an empirical matter that cannot be adjudicated here (see Ostrom 1990). Moreover, there may not be just one. It could be that different kinds of competing institutions would be preferable to a single type of institution. Again, this is an empirical issue that cannot be decided here.

Note that some institutions that would limit access to the cyber commons by Amazon et al. and have bottom-up organization would not so clearly be entrepreneurship-promoting. For example, some have proposed a "data labor union" (Imanol et al. 2018). This would limit access by Amazon et al. in the sense that such companies must engage in negotiations with a data labor union. This proposal is interesting in its own right, but it is unclear that the Kirznerian should be interested since such a union if it is not primarily focused on promoting entrepreneurial uses of that data.

Here is one way to attempt to increase entrepreneurship while giving to Amazon et al. what they are owed. We treat the cyber commons like other commons—it is either not owned or it is owned by all of us collectively, depending on one's theoretical predilections. In particular, we open up the anonymized data sets and/or metadata acquired by such companies for public use. However, to access this data, one must pay a fee. Part of this fee goes to Amazon et al. This fee allows Amazon et al. to make profits from their improvements to the cyber commons while enhancing opportunities for others to make valuable discoveries with the data that is collected and mined from Amazon et al.

How does the fee to gain access to the cyber commons get set? In a perfectly moral world, such companies would set the fee in proportion to the value they have added to the cyber commons so that they can gain a profit, but not more than they are entitled to. Figuring out this fee exactly would be quite a tall order. It would likely require some further institution—perhaps a data entrepreneurship union. Or perhaps we simply allow different companies to set their own prices for access to their data sets. Under the right conditions, this could create fair competition and therefore a market price mechanism for access to the cyber commons. Thus, even if one of two search engines is considered inferior, it could price its fee for access to its data lower than its competitor. This creates, in turn, an incentive for individuals to find creative uses for such data.

To be clear, I am instead making a moral point, not an economic one. The thought is this. By having monopoly control of the cyber commons, Amazon et al. prevent individuals from discovering new valuable uses for their data sets. If Amazon et al. are best understood as only having improved the cyber commons, they are not entitled to prevent this entrepreneurship. Given that we can further promote entrepreneurship (at least in principle) without violating the rights of Amazon et al., the Kirznerian is morally justified, *ceteris paribus*, to move toward institutions that would.

If some such institution is feasible, then the Kirznerian can say something rather nuanced and interesting about the intuition that many seem to have that we have a right to our personal digital data. While we do not have the usual property rights to our own personal digital data, we have at least some rights to our data. It is a commons. Neither Amazon nor Facebook nor Google has the right to exclude us from its use. Indeed, they have a moral duty to make it available to all of us in some form. And as a consequence, we thus have rights that it be made available to us. We are not to be excluded from the use of that data. We are not to be denied opportunities to profit from its use. Whether this view ultimately satisfies us, it provides an interesting avenue for further inquiry for the Kirznerian. It also presents an interesting lacuna in the debates over private property in our personal digital data.

NOTES

1. Of course, not all of our personal digital data is generated by anything that looks like our labor. For example, it is often useful for companies to have demographic data about us. If I fit certain demographic descriptions, I may be more likely to buy things if I am presented with certain ads. However, much of this is complicated by the fact that a lot of our demographic data is not very useful unless companies can make educated guesses about how various parts of our demographic data interact with

ever-evolving market behavior. Therefore, our demographic data is often used alongside our other data to generate predictions for ad effectiveness. The larger point here is that the real-world details about which data is being used and how matter quite a lot. The scope of my conclusions is thereby limited by the extent to which any given company is making at least some use of "labor-generated" data as opposed to some other kind.

2. We can, of course, exclude individuals from using particular instances of them. For example, in a peculiar legal scheme, someone might copyright a certain color. And every time someone uses this color, they must pay some fee back to whomever owns the copyright on that color. This could imply, of course, that the individual in question owns all instances of the abstract, universal color. According to a rather simplistic nominalist view of abstract universals, an abstract universal is just the set of all its particular concrete instances. If one holds this view, then it is perhaps possible to own an abstract universal. Because I am arguing that digital data is the kind of thing that we could own in principle, it does not bother me if, upon careful metaphysical reflection, the category of objects that cannot be owned is tinier than we would have pre-theoretically thought.

3. This worry, as far as I know, has been introduced by Glenn Weyl in a talk in 2019 for the "Future of Work" conference at the Mercatus Center at George Mason University. To my knowledge it is not yet published.

4. This is not surprising in a certain sense. It is an increasingly unpopular view among linguists and philosophers of language that the referent of a lexical item exhausts its informational content. Thus, insofar as data is just a kind of information, the problem at hand misleads us by focusing our attention on the referent.

5. By "pro tanto moral duty," I mean that there is a moral reason to not use my car without permission. But this reason could be outweighed or defeated by other moral factors. I thus reject the implausible and out-of-fashion view that those moral duties that flow from our rights are inalienable trumps across all possible contexts.

6. The interpretation of Kirzner on this count is actually a bit complicated. He sometimes appears to use Locke and neo-Lockeans like Nozick for merely expository purposes. He certainly disagrees with and finds criticisms for both. However, I think these can all be understood as in-house disputes among neo-Lockeans as opposed to criticisms from the outside. After all, Locke and Nozick, though widely discussed, did not have a monopoly on the theory of private property at the time they were writing.

7. Perhaps think about this practical-epistemic constraint on analogy with a cure for cancer. Perhaps there is presently a chemical formula that a doctor is capable of writing down and understanding that would cure cancer. We can then ask if that doctor is obligated in virtue of other individuals' entitlements to write that formula down. One reason to say that the doctor is not failing an obligation if they do not write the formula down is that the doctor does not know what the formula is. As such, we might think, for the purposes of practice, there is no obligation-inducing formula to write down.

8. It is unclear whether Kirzner would accept that, if there were other valid forms of the Creation-Entitlement Premise, the Lockean Proviso would apply to those. However, these would not, in Kirzner's view, be cases of entrepreneurial discovery.

Therefore, they may not be relevant to his defense of the justness of capitalist profits. And in any case, this exegetical point does not much matter to the inquiry at hand.

9. It might be contested that Kirzner's understanding of the relationship between justice and morality more generally is sufficiently nuanced to countenance this point. For helpful commentary, see Kirzner 2016, 124–26. However, such a move would have costs. If Kirzner availed himself of it, he might be accused of not taking justice as seriously as does his hypothetical interlocutor who wants to be convinced of the claim that pure profits in a capitalist system can be just.

10. There may, of course, be complications that arise when one considers intellectual property rights since such rights might appear, *prima facie* at least, to be property rights to types of assets. I do not mean to prejudge that issue. I do, however, suspect that if we are to believe that it's just to recognize such rights, it will not be for Kirznerian reasons. Thus, neither fans nor foes of intellectual property rights are entitled to infer much from what I say here in this chapter.

11. Perhaps there could be a limited range of data-generating search engine activities that are not themselves social or cultural, but social network websites clearly constitute social activity.

12. There are many ways one could go about emphasizing this point. Some have to do with the intrinsically social nature of economic exchange itself, the social and cultural opportunities that are provided for by market contexts, or the social nature of entrepreneurship. For various approaches to this issue, see Grube and Storr 2015. Unfortunately, it would take a whole essay to fully defend the claim that some of these features remain in an online marketplace like Amazon, so I am unable to provide this defense here. Part of my thesis thus relies on this further argument being cogent.

13. One might point out that Amazon et al. have, in a sense, already paid us back for the profits they receive in excess of what they are entitled by providing us with free services. Amazon gives impressively low-cost, fast shipping. And the use of both Google and Facebook are completely free. Some have argued that what we are owed still exceeds what we have been given in these free services (see Arrieta-Ibarra et al. 2018). This argument, of course, deserves more attention than I can give it here. One reason this argument is not entirely convincing, however, is that some data (e.g., produced by celebrities) is more valuable than others. And yet, they often enough receive the same "payment" in terms of the same free service. This suggests at the very least that the business model is not set up in order to properly compensate others. Thus, the objection must rely on the idea that it is a happy accident that we are all getting compensated to the right degree.

REFERENCES

Arrieta-Ibarra, Imanol, Leonard Goff, Diego Jiménez-Hernández, Jaron Lanier, and E. Glen Weyl. 2018. "Should We Treat Data as Labor? Moving beyond 'Free.'" *American Economic Association Papers and Proceedings* 108: 38–42.

Cheneval, Francis. 2018. "Property Rights of Personal Data and the Financing of Pensions." *Critical Review of International Social and Political Philosophy* 1–23.

Grube, Laura E., and Virgil Henry Storr, eds. 2015. *Culture and Economic Action*. Cheltenham, UK: Edward Elgar Publishing.

Kirzner, Israel M. 2016. *The Collected Works of Israel M. Kirzner: Discovery, Capitalism, and Distributive Justice*. Edited by Peter J. Boettke and Frédéric Sautet. Indianapolis: Liberty Fund.

Nolin, Jan Michael. 2019. "Data as Oil, Infrastructure or Asset? Three Metaphors of Data as Economic Value." *Journal of Information, Communication and Ethics in Society 18 (1): 28*–43.

Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.

Ostrom, Elinor. 1990. *Governing the Commons*. Cambridge, UK: Cambridge University Press.

Ritter, Jeffrey, and Anna Mayer. 2018. "Regulating Data as Property: A New Construct for Moving Forward," *Duke Law and Technology Review* 16 (1): 220–27.

Sax, Marijn. 2016. "Big Data: Finders Keepers, Losers Weepers?" *Ethics and Information Technology* 18 (1): 25–31.

Schlager, Edella, and Elinor Ostrom. 1992. "Property-Rights Regimes and Natural Resources: A Conceptual Analysis." *Land Economics* 68 (3): 249–62.