Legitimate Expectations in the Age of Innovation

Abstract: Artificial intelligence is widely expected to intensify and accelerate creative destruction, resulting in displaced workers, disrupted business models, and obsolete products. Does this violate individuals' legitimate expectations? If so, it would provide a powerful argument for blocking innovation or compensating its losers. We argue it does not. This conclusion does not depend on rejecting the doctrine of legitimate expectations. Rather, we show that the distinctive features of the market process mean that expectations about markets are rarely legitimate in the first place. As a result, the doctrine of legitimate expectations cannot justify compensating the losers of creative destruction—an implication of particular importance for current debates about AI.

Keywords: legitimate expectations; creative destruction; innovation; automation; artificial intelligence.

1. Introduction

In recent interviews, the CEO of Anthropic—the company behind the large language model (LLM) Claude—predicted that within just a few years, artificial intelligence (AI) will generate 90 percent of all software code and replace half of all entry-level white-collar jobs (O'Connor 2025; Procopio 2025). Skeptics might dismiss such claims as hype designed to inflate the value of the company. Yet both academic researchers and private sector analysts have projected equally dramatic outcomes. In one of the most influential forecasts, Carl B. Frey and Michael A. Osborne (2017: 268) estimated that "around 47% of total U.S. employment is in the high risk category. We refer to these as jobs at risk—i.e. jobs we expect could be automated relatively soon, perhaps over the next decade or two." Goldman Sachs projects that AI could automate as many as 300 million jobs worldwide (Kelly 2023). More examples are available (e.g., Brynjolfsson and McAfee 2014).

This is an example of what the economist Joseph Schumpeter (2008: 81) called "creative destruction." A bit more formally, creative destruction happens when "new innovations continually emerge and render existing technologies obsolete, new firms continually arrive to compete with existing firms, and new jobs and activities arise and replace existing jobs and activities" (Aghion et al. 2021: 1). Let's look at examples.

Innovation can render *products* obsolete. Consider how digital cameras—and later smartphones—bankrupted Kodak, how streaming services like Netflix drove Blockbuster into extinction, how Spotify and Apple Music supplanted CDs and MP3 players, how smartphones eliminated pagers and standalone GPS devices, and how electric vehicles are beginning to displace traditional gas-powered cars. Innovation can also undermine *business models*. Big-box retailers and shopping malls, once the backbone of American consumer culture, are slowly

declining as e-commerce platforms like Amazon dominate. Most unsettling of all, innovation can render certain *occupations* obsolete: industrial looms displaced artisan weavers, advanced agricultural machinery like tractor combines reduced the need for farm labor, the automobile wiped out horse-and-buggy drivers and stable hands, industrial robots replaced assembly-line workers, autonomous vehicles now threaten the livelihoods of truck and taxi drivers, and, most recently, AI is beginning to automate tasks once reserved for white-collar professionals, from computer coding to legal research.

It is easy to sympathize with those who are on the losing side of creative destruction. But does society owe them compensation? There are many ethical perspectives one might adopt to answer this question, such as the ideal of autonomy (Kogelmann 2024) or what some have called the "like cases alike" principle (Kogelmann and Carroll 2025). One might also approach the issue pragmatically: compensating the losers of creative destruction may be justified on the grounds that failing to do so risks political destabilization, such as increased support for populism (Rodrik 2018: 231, 267; Clausing 2019: 107, 112). Each of these perspectives offers a plausible framework for evaluating the ethics of creative destruction. The present paper, however, confines itself to an analysis of what the doctrine of *legitimate expectations* implies for what we owe creative destruction's losers.

What are legitimate expectations? Expectations are beliefs or predictions about the future (Brown 2017: 435). They permeate daily life: we expect a friend to keep a promise, a bowling group to meet every Thursday, or the S&P 500 to rise over the long run. Expectations are everywhere, but only some are legitimate. To call an expectation legitimate means there exists "an obligation to avoid frustrating the expectation, or to remedy the loss when the expectation *is* frustrated. It is only this subset of expectations that can, in other words, be classified as

legitimate expectations in the relevant sense" (Green 2020: 398-399). To say an expectation is legitimate thus implies two things. First, there is a reason to avoid frustrating it. This reason may not be decisive, but it is always extant. And second, when disruption of the expectation is unavoidable, we are obligated to provide compensation to its holder. Expectations that are not legitimate are "bare expectations" (Moore 2017: 231).[1]

The doctrine of legitimate expectations, as we use it, is a moral principle, but it is inspired by a legal doctrine. The legal doctrine originates in the German administrative law principle of *Vertrauensschutz* (protection of trust or reliance) and was later developed in common law systems (Forsyth 1988; Nolte 1994). In the United Kingdom, for instance, legitimate expectations arise when a public authority makes a clear promise or follows a consistent practice, thereby creating an expectation it cannot frustrate without sufficient justification or compensation (Brown 2012). Philosophers later adapted this idea as a principle of distributive justice. John Rawls (1971: 311) got there first, inspiring a substantial literature. In the philosophical sense relevant here, the doctrine holds that when individuals' beliefs about the future meet an appropriate standard, there is a *pro tanto* duty not to frustrate them, or, if they are frustrated, to provide compensation.

[1] In the philosophical and legal literature, the doctrine of legitimate expectations is a one-sided principle concerned with the moral significance of *frustrated* expectations; it does not purport to assess the moral status of expectations that are *exceeded*. Creative destruction will, of course, generate both effects: innovation will frustrate some expectations by substituting labor, while exceeding others by complementing it.

We are not sure why the doctrine of legitimate expectations is always presented and defended with this asymmetric structure. It is plausible that a case could be made for taxing and redistributing gains arising from exceeded expectations on grounds analogous to those that justify compensating frustrated ones. We thank a reviewer for raising this interesting possibility, but do not pursue it further due to space constraints and because it would detract from our main concern: what to do with creative destruction's losers.

Creative destruction no doubt frustrates expectations. The Luddites expected to work as garment weavers for the rest of their lives, but the industrial loom had other plans. The crucial question is whether such expectations are *legitimate*. If they are, the state has a reason—though not a decisive one—to guard against innovation that would undermine them. A historical example of this occurred in 1589, when Queen Elizabeth I refused to grant a patent for an automated knitting machine, because she had "too much love for my poor people who obtain their bread by the employment of knitting … to forward an invention that will tend to their ruin by depriving them of employment, and thus make them beggars" (Diamond 2019: 66). Yet given both the powerful role of creative destruction in fueling economic growth (Koyama and Rubin 2022: 9) and the profound benefits of growth (Kogelmann 2022), the reasons to restrain innovation will almost always be outweighed by competing considerations. If so, the implication is that those who lose out from creative destruction are instead owed compensation. If the Luddites' expectations are legitimate, then some form of compensation is owed after the industrial loom destroys the value of their human capital. If, however, those expectations are merely bare, then no compensation is owed.

So, were the Luddites' and similar expectations legitimate? That is our research question. Our approach is ecumenical. We do not endorse a single theory of legitimate expectations. Instead, we examine *all* relevant theories from the philosophical literature and ask: what do they imply about cases of creative destruction? The answer, in most instances, is that expectations frustrated by innovation are *not* legitimate. From this it does not follow that compensation is never owed to creative destruction's losers. There may be alternative arguments for compensation grounded in other forms of market harm. The varieties of such harms are, after all, numerous (Kogelmann and Carroll 2024). What we can say, however, is that if there *is* an

argument for compensating those harmed by creative destruction, it will seldom be supplied by the doctrine of legitimate expectations.

To anchor our discussion, it will be helpful to come up with a guiding case that we can turn to throughout the paper. This, it might be thought, is difficult, because creative destruction (as we have already noted) can manifest in many ways: it can render products obsolete, change entire business models, and eliminate jobs. Indeed, creative destruction can manifest in other ways as well. For instance, new technology like industrial robots can attenuate workers' bargaining power and hold-up opportunities, especially in capital-intensive industries (Traverso et al. 2023). What unifies these cases, however, is that they can all be understood as involving changes in relative prices. A product becomes obsolete when its market price falls below its cost of production; a business model is disrupted when alternative methods of production become relatively cheaper; an occupation is eliminated when the market wage for that type of labor falls below the opportunity costs faced by workers in the field; bargaining power and hold-up opportunities are diminished when the relative price of capital decreases. Hence, so long as our guiding case is one in which (*i*) a new innovation causes (*ii*) a shift in relative prices which (*iii*) upsets expectations, then the case should generalize, *mutatis mutandis*, to most (but perhaps not all) cases of creative destruction.

With that said, here is our guiding case:

> ALEJANDRA'S ALGORITHM. Alejandra's goal in life is to be a lawyer. She takes on significant debt to attend a law school where the median graduate makes $150,000 a year. A bit conservative, she expects to make at least $120,000 a year. While in law school, LLMs become progressively better. So much better, in fact, that they substitute many

legal services. This increases the effective supply of legal services, driving down the market price of those services. Alejandra makes $50,000 a year.

The central question is whether Alejandra's expectations are legitimate, thus giving rise to a claim for compensation. To address this, the next four sections apply the leading philosophical theories of legitimate expectations to Alejandra's case, proceeding from the most to the least plausible. None of them say Alejandra's expectations are legitimate.

2. Promising

The most uncontroversial way to generate legitimate expectations is through *promising* (Scanlon 1990; Simmons 1996: 258; Green 2020: 399; Sidgwick 1981: 269; Brown 2012: 644). Suppose Beto promises César that he will pick him up at the airport. César forms the expectation that Beto will pick him up, but Beto fails to do so. Intuitively, Beto wrongs César and owes him compensation, if only in the form of an apology. Beto owes César compensation because he upset his legitimate expectations. Beto's promise is what makes César's expectations legitimate.

Promises in general generate legitimate expectations, but there are caveats. In particular, a promise can only generate legitimate expectations among a *promisee* (the one to whom a promise is made) if the *promisor* (the one who makes the promise) has it within her power to actually bring about the state of affairs she promises to bring about. For example, if a Nashville bartender with no relationship to the music industry tells an aspiring musician, "I promise you'll be signed by a major record label if you keep playing the local music scene," then he does not generate legitimate expectations, for intuitively no compensation is owed to the musician if his big break never materializes.

What underpins this intuition? Those who theorize about promising argue that certain *uptake conditions* must be in place for a promisor to successfully make a promise (Liberto 2018). For example, you cannot make a promise if someone is pointing a gun at your head demanding you do so, nor can you make a promise if under the influence of a mind-altering drug. Likewise, promisors are unable to genuinely make a promise to φ if it is not within their power to φ. Because the Nashville bartender cannot control whether the musician gets a record deal, he cannot make that promise. Because the bartender fails to make a promise, and because promises are what underwrite legitimate expectations (according to the current theory), the musician's expectations are not legitimate.

This leads us to the following theory of legitimate expectations:

PROMISE: Individual *i*'s expectations that outcome *O* will obtain are legitimate if and only if an agent who has the capacity to ensure *O* will obtain promises *i* that *O* will obtain.

PROMISE explains both the case that opened this section regarding Beto and César, as well as the Nashville bartender case. In the Beto-César case, Beto has it within his power to pick César up at the airport and promises to do so. Therefore, César's expectations are legitimate, and he is owed compensation if Beto fails to pick him up. In the other case, the bartender does not have it within his power to ensure the musician receives a record deal, but promises him anyways. Therefore, the musician's expectations are not legitimate, and he is owed no compensation if he fails to land a major record deal.

Does PROMISE say Alejandra's expectations are legitimate? For them to be, someone must have first promised her that she would make at least $120,000 a year if she attends law school. Most cases of creative destruction, we suspect, will not satisfy this condition. Very rarely do individuals literally *promise* others a certain salary (or that they will always work in a specific

occupation, or that a specific business model will always be viable, etc.). Note, *suggesting* one will earn a specific salary is not the same as *promising* it. A guidance counselor who says, "Most people who go to the law school you've been accepted to make $150,000 a year, so you'll probably make at least $120,000," does not make a promise, but rather a suggestion based on available evidence. Perhaps this suggestion is sufficient to generate legitimate expectations—more on this in the next section—but it is not according to PROMISE.

But suppose someone does promise Alejandra that she will make at least $120,000 a year if she attends law school. A promise like this might come from her parents, friends, teachers and professors, guidance counselors and career advisors, or perhaps even the law school itself. But this is not yet sufficient to generate legitimate expectations according to PROMISE, for it must be the case that these individuals have it within their power to bring about the relevant state of affairs. That is, they must have it within their power to ensure the market price for Alejandra's labor is $120,000.

And this brings us to the main problem with applying PROMISE to cases of creative destruction. The expectations frustrated by creative destruction, we argued in the prior section, can almost always be construed as expectations about future market prices (e.g., an expectation that one's product will never be rendered obsolete can be construed as an expectation about the future market price of that product). But almost never will a single individual or entity be able to control a price. Market prices result from the uncoordinated actions of millions of buyers who try to get the lowest price possible and millions of sellers who try to get the highest price possible (Hayek 1945). These decentralized interactions typically render all participants *price takers*: individuals who cannot influence prices and must treat them as externally fixed (Mankiw 2015: 66). By definition, a price taker cannot generate legitimate expectations for Alejandra, for they

do not have it within their power to ensure she makes $120,000, even if they promise her she will make that.[2]

When markets are imperfect, not everyone is a price taker. This occurs when someone possesses significant market power. For instance, a buyer or seller who accounts for a large share of market transactions—say, purchasing or selling 20 percent of a given good—may be able to

---

[2] One might object that our reasoning here amounts to a *reductio* of our definition of PROMISE, showing it to be implausibly demanding. More specifically, suppose *A* signs a contract with *B* to supply a widget at a fixed price. As we have just argued, however, *A* cannot control the prices of all the components that go into producing the widget. If control over performance is a necessary condition for making a promise, then *A* cannot genuinely promise to supply the widget at that price, for the price of its components may rise steeply outside of *A*'s control. If so, the contract would appear void *ab initio*—an absurd result, since it would imply that ordinary commercial exchange rests on a mistake, for most contracts are incomplete in that they cannot address all contingencies (Williamson 1979; Vatiero 2021).

This objection has force only if contracts are understood as promises. Some legal theorists embrace this view (Fried 1981). Others, however, reject the identification precisely because it generates objections of the sort just described. As Daniel Markovits (2021: §1.1) puts it, "a central challenge for the promise conception [of contracts] is the divergence of well-established legal rules from the rules of promissory morality." Take another example: one-sided promises—for example, a promise to meet a friend for coffee—can clearly be *bona fide* promises, yet they are typically not enforceable as contracts. The divergence between promising and contracting suggests the two are distinct phenomena.

Once this is recognized, alternative theories of contract become available (Markovits 2021: §1). On our preferred view, contracts are institutional devices designed to maximize economic productivity (Schwartz and Scott 2003). They take on promissory features only insofar as those features serve the end of economic productivity. When a feature characteristic of promising—such as strict control over performance—would impede productive exchange, it is appropriately set aside.

Returning to the case at hand, *A*'s contractual agreement to sell *B* a widget at a fixed price may fail to qualify as a genuine promise (because *A* cannot control the outcome), but it can nonetheless constitute a perfectly valid legal contract. If *A* fails to deliver, *B* is owed compensation *not* because her legitimate expectations are upset (no genuine promise has been made), but because an economically efficient contract law demands it.

influence prices through the scale of their activity (Mankiw 2015: 150). In the current case, if one firm purchases a substantial portion of legal services in a particular market, it may be able to affect the prevailing wage by adjusting its willingness to pay. Such cases are rare. But even when they arise, they still fall short of meeting the conditions required by PROMISE. *Influencing* a price is not the same as *determining* it. Even a dominant market actor is unlikely to have the capacity to ensure Alejandra will earn exactly $120,000; at most, they can exert modest pressure on the wage level. This limited capacity to affect outcomes does not suffice to ground legitimate expectations according to PROMISE.

Who has the capacity to determine the market price of Alejandra's labor? The state can. The state can put a price floor on legal services. The likely result of this is unemployment (i.e., a surplus of legal services), because the state is requiring consumers to pay above the market price for a good (Mankiw 2015: 117). Instead of a price floor, the state can subsidize legal services, paying the difference between what consumers are willing to pay and $120,000. Or, if the state goes with the price floor option and causes unemployment, it can simply employ all the lawyers who are no longer employable at $120,000.

Though the state has the capacity to ensure Alejandra makes $120,000 if she goes to law school, it is incredibly unlikely the state would ever promise this. We need to be very clear about what a promise from the state looks like. A promise from the state is not the same as a promise from someone affiliated with the state. If a congressman promises Alejandra that she will make $120,000 a year if she goes to law school, this is not sufficient to generate a legitimate expectation, because the congressman, on his own, lacks the capacity to ensure this. If the President of the United States promises Alejandra that she will make $120,000 a year if she goes

to law school, this will not generate legitimate expectations because the president, on his own, lacks the capacity to ensure this.

In the U.S., the entity that has the capacity to secure Alejandra a certain salary is a majority in the House of Representatives, a supermajority in the Senate (due to the filibuster), and the President acting in tantamount. If this complex corporate agent promises Alejandra that she will make $120,000 a year acting as a lawyer, then Alejandra's expectations are legitimate. But this would be very strange indeed. Essentially, the U.S. government would have to pass a statute promising Alejandra (or perhaps all lawyers) this. This is logically possible, but describes very few (if any) cases. We can think of none. PROMISE thus does not say Alejandra's (and similar) expectations are legitimate.

3. Encouragement

Here's an example of legitimate expectations being violated. Dulce works for a company and desperately wants a promotion. Her boss constantly says things like "If you keep coming in early and staying late, you'll probably get promoted," and "We are very impressed with your performance; if you keep up the good work, you're almost certain to get a promotion." Dulce forms the expectation that if she achieves the performance markers suggested by her boss then she will be promoted. From the sample conversations provided it should be clear that Dulce's boss promises her nothing. And yet, if the promotion never comes after Dulce achieves all the relevant performance markers her boss suggests, then intuitively she has been wronged, and is owed compensation (in particular, she is owed the promotion). Dulce's expectations are legitimate.

But Dulce's boss promised her nothing, so what underwrites the legitimacy of her expectations? A plausible answer is that, though her boss promises her nothing, he was responsible for creating the expectation that she will receive a promotion if she achieves certain performance markers. Put another way, Dulce's boss *encouraged* the expectation. Encouragement, just like promising, can underwrite the legitimacy of expectations. This is the theory of legitimate expectations that Alexander Brown develops for the administrative state specifically. He writes:

> … expectations about what governmental administrative agents or agencies will do or not do in the future are legitimate if, and only if, those agents or agencies were responsible for creating the expectations, after they had been given or had assumed a role responsibility, competence, or discretion over the relevant policies and measures (Brown 2017: 445).

Beyond Brown's theory of legitimate expectations for the administrative state, other philosophers argue that encouragement or inducement is what renders expectations legitimate (Rawls 1971: 311; Hsieh 2000: 102; Quigley 2023: 710;).

As with promises, encouragement can only render expectations legitimate if the encourager has it within her power to bring about the relevant state of affairs. This is explicit in Brown's account, as illustrated in the quote above. To see why, suppose it is not Dulce's boss, but rather a janitor at her firm, that continually tells Dulce "If you keep coming in early and staying late, you'll probably get promoted," and so on. The janitor's words of encouragement may lead Dulce to form an expectation that she will be promoted, but intuitively she is owed no compensation if the promotion never comes. The reason why is because the person encouraging

Dulce's expectations does not have it within his power to ensure her expectations come to fruition.

We can now state a second theory of legitimate expectations:

> ENCOURAGEMENT: Individual *i*'s expectations that outcome *O* will obtain are legitimate if and only if an agent who has the capacity to ensure *O* will obtain encourages *i* to believe that *O* will obtain.

ENCOURAGEMENT explains our intuitions in the two Dulce cases. When her boss encourages the expectation that she will be promoted, the expectation is legitimate, for her boss has the capacity to make the expectation come true. But when the janitor encourages the expectation that she will be promoted, the expectation is not legitimate, for the janitor lacks the capacity to ensure the expectation comes true.

Are Alejandra's expectations legitimate, according to ENCOURAGEMENT? We noted in the section above that in very few cases do individuals literally *promise* others that they will make a certain salary. Yet, many *encourage* such beliefs. For instance, when Alejandra's guidance counselor says, "Most people who go to the law school you've been accepted to make $150,000 a year, so you'll probably make at least $120,000," she encourages an expectation of a future salary. Moreover, if Alejandra's friend says, "My sister went to the same law school you've been accepted to and she makes $120,000 a year," she too encourages an expectation of a future salary. Such expectations are only legitimate, though, if those who encourage them have it within their power to bring about the relevant state of affairs. And as we noted in the section above, in almost no cases will individuals or entities other than the state have it within their power to ensure a specific market price obtains. Guidance counselors, parents, and friends might all

encourage Alejandra to believe she will make a certain salary if she goes to law school, but they do not have the power to make this come true.

There is one entity, we argued in the prior section, that does have the power to determine prices: the state. Though the state has the capacity to determine prices we argued that rarely (if ever) does it promise certain prices. But though the state rarely promises, it often encourages. This encouragement can generate legitimate expectations, some argue. Consider an example. Travis Quigley (2023) argues in defense of private land use regulations (i.e., zoning) because homeowners have legitimate expectations that their neighborhoods will remain relatively unchanged over a long period of time.[3] These expectations arise because the state encourages them via subsidizing 30-year mortgages (Quigley 2023: 712). Because the state encourages expectations that neighborhoods will remain stable, and because the state has the power to ensure this, the expectations are legitimate. Therefore, if neighborhoods change quickly and significantly, then compensation is owed, according to both Quigley (2023: 719) and ENCOURAGEMENT.

The same reasoning can be applied to ALEJANDRA'S ALGORITHM. The state encourages Alejandra to attend law school by subsidizing student loans. In response, Alejandra attends law school and forms expectations about her future career as a lawyer. Because the state encouraged her expectations, and because the state has the power to ensure they come true, Alejandra's expectations are legitimate. According to ENCOURAGEMENT, she is owed compensation when new AI technology dramatically reduces her wage.

[3] For criticism of Quigley's argument, see Kogelmann (2025).

We need to be careful here. The above cases conflate two things: the state encouraging individuals to pursue an *action* and the state encouraging individuals to form an *expectation*. To encourage an action φ is to give someone a reason to do it; to encourage an expectation that outcome *O* will obtain is to give them a reason to believe it. For ENCOURAGEMENT to say an expectation is legitimate, the state must do the latter, not the former. But in the cases above the state only does the former. By subsidizing 30-year mortgages, the state gives individuals a reason to buy a home. But how does subsidizing 30-year mortgages give individuals reason to believe their neighborhoods will retain their character in the long run? By subsidizing student loans, the state gives Alejandra a reason to attend college and postgraduate school. But how does subsidizing student loans give her a reason to believe she will make at least $120,000 over the course of her career?

One might reply: when you give someone a reason to φ, it's quite natural that they also acquire a corresponding belief that, should they φ, a specific outcome *O* will obtain, even if they are not given an explicit reason to believe *O* will obtain. Hence, by encouraging someone to φ, we also in some sense encourage them to believe *O* will obtain. This leads to the following revision:

> DERIVATIVE ENCOURAGEMENT: Individual *i*'s expectations that outcome *O* will obtain are legitimate if and only if someone encourages *i* to φ and *i* believes *O* is the likely result of *i*'s φ-ing.

There are problems with this view. Alejandra, suppose, is plagued by hubris. She believes that if she were to attend law school she would make at least $5,000,000 a year. The state subsidizes student loans, encouraging Alejandra to attend law school. She does, and then forms the hubristic

expectation. But no one thinks Alejandra deserves compensation if she makes less than $5,000,000 a year.

Perhaps something like DERIVATIVE ENCOURAGEMENT will work if we limit the kinds of expectations individuals can form around actions they are encouraged to do. One attempt to do this looks at the intentions of the encourager. Typically, when someone encourages another to perform an action, they do so with the intention that a specific outcome will be realized. If the individual who has been encouraged comes to expect this intended outcome, then such expectations are legitimate. More formally:

> INTENDED ENCOURAGEMENT: Individual *i*'s expectations that outcome *O* will obtain are legitimate if and only if individual or entity *j* encourages *i* to φ with the intention that *O* obtains, and *i* believes *O* is the likely result of *i*'s φ-ing.

This revised principle also faces problems, especially when applied to actions encouraged by the state or other corporate entities. The worry is that it's not clear how to make sense of intentions in these cases.

In terms of the state, one can point to the intentions of those who pass a particular bill, but there will often be many distinct and conflicting intentions behind it. In the classic bootleggers and Baptists example, alcohol was prohibited due to a coalition of bootleggers who wanted a monopoly on alcohol sales and Baptists who wanted to enforce their religious creed (Smith and Yandle 2014). What was the state's intention, in this case? Similar issues will arise with ALEJANDRA'S ALGORITHM. When the legislature passes a bill subsidizing higher education many intentions will likely be present: to ensure citizens earn good incomes, to grant rents to the higher education industry, to foster human capital to produce a more innovative economy, to

encourage marriage, to reduce crime, to increase long-term tax revenues, to gain favor with younger voters, etc. Given these complexities, it's hard to know how to actually apply INTENDED ENCOURAGEMENT in cases of corporate actors.

In sum, ENCOURAGEMENT is a plausible theory of legitimate expectations, but probably renders few cases of expectations frustrated by creative destruction legitimate, because only the state can ensure specific market prices obtain, but the state rarely encourages its citizens to form *expectations* about specific market prices. The state often encourages individuals to perform certain *actions*, but this is distinct from encouraging expectations. Various attempts were made to develop a theory of legitimate expectations from the state encouragement of actions, but these attempts failed.

4. Social Norms

A third account says legitimate expectations are generated by *social norms* or *social practices* (Moore 2017; Green 2020; Melenovsky 2021). Fergus Green (2020: 402) motivates this view by comparing two cases, one of which is taken from A. John Simmons (1996: 258). In the first case, Edgar and Felipe are good friends and have dinner with each other every Friday evening, taking turns as hosts. One Friday, Edgar prepares dinner; without warning or notice, Felipe does not show up. The intuition is that Felipe wrongs Edgar and owes him compensation, if only in the form of an apology. In the second case, drawn from Simmons, Kant takes his daily walk through the city of Königsberg at precisely the same time each day—so reliably, in fact, that local housewives use his passing as a cue to set their watches. Should Kant one day fail to

take his walk, it seems intuitive that he owes no compensation to the housewives, even though they are left uncertain of the time.

Legitimate expectations are present in the Edgar-Felipe case, but not the Kant one. Why? One plausible answer is that a *social norm* is violated in the former, but not the latter. Moore (2017: 231) writes: "A legitimate expectation … arises from the operation of rules, practices, policies, norms and conventions." We can state this view as follows:

> NORMS: Individual *i*'s expectations that outcome *O* will obtain are legitimate if and only if there exists a social norm that, if adhered to, will produce *O*.

In the Edgar–Felipe case, it's plausible to think a social norm is present: the two are to have dinner every Friday night, alternating as hosts. This norm grounds Edgar's expectations, rendering them legitimate. By contrast, Kant's failure to take his usual walk does not breach any social norm. Although there was a consistent pattern of behavior, a pattern of behavior is not the same as a norm. The housewives' expectations are not legitimate.

How can we be sure a norm is present in the Edgar-Felipe case but not the Kant one? To be confident of this judgment, and to apply NORMS more generally, we need a theory that identifies when social norms exist. The most well-established theory from the social sciences is Cristina Bicchieri's (2006; 2017). She defines social norms as follows:

> A social norm is a rule of behavior such that individuals prefer to conform to it on the condition that they believe that (*a*) most people in their reference network conform to it (empirical expectations), and (*b*) that most people in their reference network believe they ought to conform to it (normative expectations) (Bicchieri 2017: 35).

This definition requires explanation. To begin, a *reference network* is the group to whom a social norm applies. Some norms operate in small towns (everyone attends their neighbor's funeral), others in big cities (don't make eye contact on the subway), and still others across entire countries (shake with your right hand).

According to Bicchieri's theory, for a social norm to exist within a particular reference network, three conditions must be met. First, there must be a recognizable behavioral rule, such as "don't litter" or "greet others with a handshake." Second, there must be *empirical expectations*: individuals believe that others in their reference network generally conform to the rule. In the case of "don't litter," for example, most people in the network must believe that others refrain from littering. Third, there must be *normative expectations*: individuals believe that others in the network think the rule ought to be followed. That is, most people must believe that others should not litter. Only when all three elements are in place does a social norm exist, according to Bicchieri.

In the Edgar-Felipe case, all three of Bicchieri's conditions are satisfied. The reference network consists solely of Edgar and Felipe. First, there is a behavioral rule: the two have dinner together every Friday, taking turns to host. Second, empirical expectations are present—Edgar believes that Felipe will follow the rule, and Felipe believes that Edgar will follow the rule. Third, normative expectations are also met—each believes the other ought to follow the norm. Jointly, these conditions establish the presence of a social norm, thereby rendering Edgar's expectations legitimate.

The Bicchieri framework is more difficult to apply to the Kant case, suggesting no social norm is present. To begin, it's unclear what the relevant behavioral rule is. It cannot be "Kant must walk through Königsberg at the same time every day," for many think rules and norms

must be *general*, in that they apply to classes of people, not individuals (Rawls 1971: 131). If we instead propose the rule "The people of Königsberg must walk through the city at the same time every day," then both empirical and normative expectations are not met. Empirical expectations are not met because only Kant follows the rule. Normative expectations are likewise absent, as it is implausible that residents believe others ought to take the same daily walk. The housewives' expectations are not legitimate.

What does NORMS say about Alejandra's expectations? It says they are legitimate only if there exists a social norm stating that lawyers must earn at least $120,000 annually. To determine whether such a norm exists, we must identify the relevant reference network and then assess whether the three conditions are satisfied.[4]

What would the reference network be in this case? Plausibly, it would consist of participants in a specific market for legal services. For example, the legal services market in

[4] One might object that these steps are unnecessary. Bicchieri (2006: 11) characterizes social norms as equilibria, and in equilibrium agents lack incentives to unilaterally alter their behavior. If this is right, then where a genuine norm exists, Alejandra's expectations could not be frustrated, since no other actor would have an incentive to engage in conduct that undermines them.

It is true that Bicchieri describes social norms as equilibria, and it is true that agents lack incentives to unilaterally deviate from an equilibrium. Nonetheless, the equilibria Bicchieri analyzes are fragile. They depend essentially on empirical and normative expectations, and when those expectations change, the equilibrium dissolves. Such shifts can occur for many reasons, and much of Bicchieri's work is devoted precisely to explaining how norms emerge, erode, and disappear as expectations change (e.g., Bicchieri 2002; Bicchieri and Funcke 2018). Indeed, she has worked for Oxfam to change social norms in developing countries by manipulating expectations (Bicchieri 2017). A social norm may therefore exist at one time even though expectations are later frustrated due to the norm being violated. In such cases, expectations are thwarted not because no norm ever existed, but because the conditions sustaining the norm no longer obtain.

Dallas–Fort Worth might constitute one reference network, while the market in Los Angeles would constitute another. Each of these regional markets could plausibly sustain its own norms about lawyer compensation. To further flesh out our guiding case, suppose Alejandra attends the UCLA School of Law and plans to enter the Los Angeles legal market.

Now to the three criteria. First, is a behavioral rule present? Plausibly, yes. The relevant rule would be something like: "Lawyers must earn at least $120,000 per year." Is this rule sufficiently general to count as a genuine social norm? Recall our earlier argument: the rule "Kant must walk through Königsberg at the same time every day" fails the generality test because it applies to a single individual. By contrast, "Lawyers must earn at least $120,000 per year," is sufficiently general, even though it does not apply to everyone. Generality is not the same as universality (Rawls 1971: 131-132). It only requires that the rule apply to a group in virtue of some generalizable characteristic—in this case, being a lawyer.

Second, are empirical expectations present? In Bicchieri's framework, this requires that individuals within the relevant reference network believe that others in the network conform to the rule. Applied to Alejandra's case, this would mean that actors in the relevant legal market—prospective lawyers, employers, clients, and so on in Los Angeles—believe that lawyers in that market generally earn at least $120,000 per year. Empirical expectations are not about what people think *ought* to happen, but what they believe *actually* happens. It's not difficult to imagine this condition being satisfied, especially if salaries near or above that threshold are common in the region.

Difficulties arise when we consider the third condition: normative expectations. For it to count as a social norm in the Los Angeles legal market that lawyers must earn at least $120,000 per year, individuals in the market must not only believe that lawyers *do* earn that amount, but

also that lawyers *should* earn that amount. This, however, is implausible. In large and heterogeneous markets, we should expect a wide range of views regarding just prices (Geisst 2023). Some may believe that lawyers are underpaid, others may be indifferent, and still others (in particular, their clients) may believe they are overpaid. This variation in normative judgment undermines the uniformity required for a social norm to exist. So long as there is significant heterogeneity in what people believe lawyers *ought* to be paid, the shared normative expectations condition is not met. With no norm present, Alejandra's expectations are not legitimate.

This brings us to a broader challenge with applying NORMS to expectations frustrated by creative destruction. To apply it, there must be social norms specifying what the prices of particular goods or services must be. There is nothing conceptually incoherent about this, but we should expect such norms to be exceptionally rare, for norms require normative expectations, and thus a significant degree of normative agreement. All else equal, normative agreement is increasingly difficult to achieve as the size and diversity of a group grows. This is why we often find *thick* norms—those involving substantive moral or cultural commitments—in very small, relatively homogenous communities (e.g., that everyone in a small town must observe a particular religion). By contrast, large and diverse communities, such as an entire country or continent, sustain only *thin* norms (e.g., minimal behavioral expectations like greeting one another with a handshake).

In today's globalized world, markets are large, encompassing buyers and sellers with widely differing backgrounds, values, and interests. A norm specifying a particular market price would be thick, since it presumes a high degree of normative consensus about what constitutes a just price. But in large markets, such normative consensus will almost never be found. No normative consensus surrounding just prices means no normative expectations, which means no

social norms about prices. For this reason, NORMS will almost never say that expectations frustrated by creative destruction are legitimate.

5. Patterns

The final and most controversial theory of legitimate expectations grounds them in patterns or regularities. According to this view, when a certain event has occurred with sufficient regularity in the past, individuals are epistemically justified in expecting it to continue in the future. Precisely because their expectations are epistemically justified, they are also legitimate. While many scholars who work on this topic think epistemic justification is a *necessary* condition for expectations to be legitimate (Breakey 2022: 2808; Quigley 2023: 712; Meyer and Truccone 2025: 5), the current view says epistemic justification is *sufficient* for them to be legitimate. Robert Goodin writes:

> Nothing is unreasonable, statistically, in predicting that a settled state of affairs will persist in much the same form. Induction is not *that* unwarranted. The future is indeed only marginally different from the present, very much more often than not … we compensate to restore the status quo ante only if it has been upset either (a) in ways people had no reason to expect or (b) in ways people had reason to expect not to occur (Goodin 1991: 153).

Others advance a similar view. For instance, Lukas H. Meyer and Santiago Truccone-Borgogno (2022: 704) say that expectations are legitimate when "some past behaviour or circumstance gives reasons to believe that the future will be akin to the past in virtue of which the expectation is formed." We can state this view as follows:

PATTERNS: Individual *i*'s expectations that outcome *O* will obtain are legitimate if and only if *i*'s expectations are epistemically justified; regular occurrence of *O* in the past is sufficient to epistemically justify *i*'s expectations.

PATTERNS is controversial. Consider the two cases from the prior section. While it is intuitive that Felipe owes Edgar some kind of compensation for violating their shared dinner norm, it is counterintuitive that Kant owes the housewives of Königsberg anything for skipping his daily walk. Yet a plausible application of PATTERNS says he does. Because Kant walked at the same time every day for a year, the housewives were epistemically justified to expect him to continue doing so. When he failed to appear, their expectations were upset. According to PATTERNS, those expectations are legitimate, so compensation is owed.

One might think PATTERNS supports the conclusion that Alejandra's expectations are legitimate. Her expectation that she will earn at least $120,000 is, let's assume, grounded in a body of historical data: for the past fifty years, lawyers who graduated from her law school consistently earned at least that amount. Given the long-standing nature of the pattern, Alejandra is epistemically justified in expecting the pattern to continue. By PATTERNS, her expectations are therefore legitimate.

We have two objections to this conclusion. The first concerns whether Alejandra's expectations are being judged against the "right" pattern.[5] In the paragraph above, the pattern that renders Alejandra's expectations epistemically justified and hence legitimate is the fact that, for the past fifty years, lawyers who graduate from her school consistently earned at least $120,000 a year. But there is another relevant pattern that points to the opposite conclusion.

[5] We thank an anonymous reviewer for suggesting this argument.

When looking at the broad history of capitalism, creative destruction is a relatively common phenomenon (Greenspan and Wooldridge 2018). Schumpeter (2008: 84) called it a *perennial* gale for a reason. More generally, markets are and have always been fraught with uncertainty (Knight 2014). If this is the relevant pattern that we judge Alejandra's expectations by, then her expectations are not epistemically justified—for they completely ignore the historical pattern of creative destruction—and therefore are not legitimate. We are unsure which pattern is the "correct" one by which to judge the legitimacy of Alejandra's expectations (defenders of PATTERNS are not clear on this point). But on at least one plausible way of applying PATTERNS, Alejandra's expectations are not legitimate

But to grant everything we can to our interlocutors, assume now that the "right" pattern to judge Alejandra's expectations by is the fact that lawyers who have graduated from her school for the past fifty years have made at least $120,000 a year. This brings us to our second objection. Here we do not challenge the core idea that if expectations are epistemically justified, they are also legitimate. We instead add some nuance to the second component of PATTERNS, which says regular occurrence of *O* in the past is sufficient to render expectations that *O* will occur again in the future epistemically justified. In some cases, inferences about the future based on longstanding patterns in the past are *not* epistemically justified, we argue. To see why, consider a case.

> KANT'S LUCKY WALK: Kant is famously regular in his habits, taking the same walk through Königsberg every day at precisely the same hour. On his route, he must cross ten drawbridges. The drawbridge operators, however, are notoriously unreliable: they sometimes fail to show up, demand bribes, or mistime the lowering of the bridges. Yet over the past year, Kant has been extremely lucky: each day he has arrived at the bridges

> just as they were lowered, allowing him to complete his walk uninterrupted. The townspeople, especially the housewives who watch for his punctual appearances, are well aware of the erratic behavior of the drawbridge operators.

Because we want to work within the PATTERNS framework, we grant for the sake of argument that the housewives' expectations are legitimate in the ordinary Kant case introduced in the prior section. The question is whether they are legitimate in KANT'S LUCKY WALK as well. Our intuition is no. Even if the housewives are owed compensation in the ordinary case, they are not in this one.

Why are the expectations in KANT'S LUCKY WALK not legitimate, despite the past regularity of his daily routine? Because the pattern the housewives observe is *fragile*. A pattern is fragile when it depends on the precise alignment of numerous underlying causal factors—like ten unreliable drawbridges—and when the failure of even one of these factors can disrupt the entire pattern. What we call fragile patterns are often referred to by others as *complex systems* (Page 2011). In complex systems, many interdependent components work together to produce an outcome. A minor disturbance of just one component can throw the entire system off course. When dealing with these kinds of systems, it is not epistemically reasonable to infer that the future will mirror the past, since the system is highly sensitive to small changes. If one knows a pattern is produced by a fragile system, then one should expect that pattern to be disrupted at any time regardless of how it behaved in the past. This is why the expectations in KANT'S LUCKY WALK are not legitimate, even though similar expectations might be legitimate in the ordinary Kant case.

We do not challenge the core principle of PATTERNS, which says epistemically justified expectations are legitimate. Our concern lies with the assertation that observation of past patterns

is sufficient to render expectations epistemically justified. This is true only in some cases. It fails when the patterns in question are produced by fragile or complex systems. The relevant question, then, is whether Alejandra's expectations are about a fragile, complex system or whether they are about a relatively simple system.

Alejandra's expectations are about market outcomes, and markets are paradigmatic examples of complex systems (Foster 2005; Kuhlmann 2014; Vargo et al. 2017). They consist of numerous interacting agents—consumers, firms, investors, regulators, etc.—each pursuing their own objectives based on limited and often asymmetric information. The aggregate behavior of the market emerges from their decentralized decisions and cannot be fully explained by analyzing individuals in isolation. Markets exhibit nonlinear dynamics, where small events can have disproportionately large effects, such as speculative bubbles or crashes. Feedback loops—both positive, as in herd behavior, and negative, as in corrective price mechanisms—further complicate predictions. Agents within the system constantly adapt to new information and shifting incentives, leading to path-dependent outcomes and constant change. These features are hallmarks of complexity, making markets inherently difficult to predict and prone to sudden, sometimes chaotic, shifts.

All this applies to ALEJANDRA'S ALGORITHM. The legal labor market is a complex system. It is shaped by the decentralized interactions of thousands of actors—lawyers, law firms, clients, law schools, policymakers, regulators, developers of legal AI, etc.—each acting with incomplete information and responding to shifting incentives. Alejandra's expectations that lawyers will continue to earn at least $120,000 are based on a historical pattern, but this pattern emerged from a delicate balance of contingent factors: the demand for legal services, the supply of legal services, the structure of professional licensing, wage norms in elite firms, the exclusion

of non-lawyers from the provision of many legal services, and more. A small change—for instance, the bar association allowing non-lawyers to perform more routine legal tasks—can radically alter this balance.

Because the legal market's prior behavior depended on a fragile constellation of interdependent factors, even a modest change to one of them can upend existing patterns. As such, Alejandra's expectations about the future market price of legal services are not epistemically justified. One cannot, we have argued, form epistemically justified expectations about the future behavior of fragile systems merely by observing past performance. But if epistemic justification is lacking then expectations are not legitimate, according to PATTERNS. Alejandra's case illustrates why PATTERNS will almost never say expectations frustrated by creative destruction are legitimate: in complex systems such as markets, past regularities cannot epistemically justify predictions about the future.

6. Conclusion

Artificial intelligence has become the focal point of contemporary anxieties about creative destruction. Unlike earlier technological shifts, AI appears capable of simultaneously affecting a wide range of occupations, including those previously thought insulated from automation. Its speed, scale, and generality make the prospect of disruption feel unprecedented. As a result, many commentators agree that those who lose from AI-driven creative destruction should be compensated. Indeed, many of the AI creators and companies *themselves* are calling for robust compensatory and redistributive policy responses, such as universal basic income (Altman 2014; Amodei 2024).

But what remains far more contested is *why* such compensation is owed. Answering this "why" is not merely an intellectual exercise (although, as philosophers, we do think it is intellectually interesting). Different normative justifications for compensation imply different forms it should take. Compensation grounded in legitimate expectations, for instance, should be backward-looking and individualized, aiming to make particular individuals whole for specific frustrated expectations.

By contrast, compensation grounded in considerations of distributive justice, risk-sharing, autonomy, or political stability will point toward very different institutional responses. A distributive-justice rationale is likely to support broad, rule-based tax-and-transfer schemes aimed at correcting inequalities generated by innovation, rather than individualized remedies to make individuals whole. A risk-sharing or insurance-based rationale instead treats innovation-induced losses as social risks to be pooled *ex ante*, favoring mechanisms such as wage insurance or social insurance programs rather than compensation tied to particular disappointed expectations. An autonomy-based justification emphasizes individuals' capacity to revise life plans in the face of technological change, suggesting policies focused on retraining, mobility, and the preservation of meaningful choice rather than income replacement alone. Finally, arguments grounded in political stability may justify compensation that is targeted, temporary, or strategically deployed to mitigate social unrest and the turn toward populism, even if no individual injustice has occurred.

Thus, there is a need for a normative theory to guide how compensation should be pursued for the losers of AI-induced creative destruction. The doctrine of legitimate expectations is a plausible candidate, as some have recently suggested (Belic 2024). But this paper has argued that the doctrine of legitimate expectations *cannot* justify compensating those who will lose from

AI-induced creative destruction. The reason is not that AI-induced disruptions are morally insignificant, but that the kinds of expectations most commonly frustrated by innovation—including expectations about wages, occupational stability, and the continued viability of particular skills—are rarely legitimate in the relevant sense. Across the leading philosophical accounts of legitimate expectations—PROMISE, ENCOURAGEMENT, NORMS, and PATTERNS—market-based expectations typically fail to meet the conditions required to generate *pro tanto* claims against frustration.

Importantly, this conclusion should *not* be read as a rejection of compensation, redistribution, or regulation in the age of AI. It is instead a claim about *where such arguments must come from*. Concerns about autonomy, fairness, populism and political unrest, or the distribution of social risk may well support robust policy responses to AI-driven disruption. What our argument shows is that the doctrine of legitimate expectations cannot bear this justificatory weight. But we believe something else can, and more philosophical work must be done to determine what can bear this weight in order to figure out what the appropriate policy response to AI should be.

Works Cited

Aghion, Philippe, Céline Antonin, and Simon Bunel. 2021. *The Power of Creative Destruction*. Harvard University Press.

Amodei, Dario. 2024. "Machines of Loving Grace." URL: https://www.darioamodei.com/essay/machines-of-loving-grace#5-work-and-meaning.

Altman, Sam. 2014. "Technology and Wealth Inequality." URL: https://blog.samaltman.com/technology-and-wealth-inequality.

Belic, Jelena. 2024. "Institutions, Automation, and Legitimate Expectations." *Journal of Ethics* 28: 505-525.

Bicchieri, Cristina. 2002. "Covenants without Swords: Group Identity, Norms, and Communication in Social Dilemmas." *Rationality and Society* 14: 192-228.

Bicchieri, Cristina. 2006. *The Grammar of Society*. Cambridge University Press.

Bicchieri, Cristina. 2017. *Norms in the Wild*. Oxford University Press.

Bicchieri, Cristina and Alexander Funcke. 2018. "Norm Change: Trendsetters and Social Structure." *Social Research: An International Quarterly* 85: 1-21.

Breakey, Hugh. 2022. "When Normal is Normative: The Ethical Significance of Conforming to Reasonable Expectations." *Philosophical Studies* 179: 2797-2821.

Brown, Alexander. 2012. "Rawls, Buchanan and the Legal Doctrine of Legitimate Expectations." *Social Theory and Practice* 38: 617-644.

Brown, Alexander. 2017. "A Theory of Legitimate Expectations." *Journal of Political Philosophy* 25: 435-460.

Brynjolfsson, Erik and Andrew McAfee. 2014. *The Second Machine Age*. W.W. Norton & Company.

Clausing, Kimberly. 2019. *Open: The Progressive Case for Free Trade, Immigration, and Global Capital*. Harvard University Press.

Diamond, Arthur M., Jr. 2019. *Openness to Creative Destruction*. Oxford University Press.

Forsyth, C.F. 1988. "The Provenance and Protection of Legitimate Expectations." *Cambridge Law Journal* 47: 238-60.

Foster, John. 2005. "From Simplistic to Complex Systems in Economics." *Cambridge Journal of Economics* 29: 873-892.

Frey, Carl B. and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible are Jobs to Computerization?" *Technological Forecasting and Social Change* 114: 254-280.

Fried, Charles. 1981. *Contract as Promise: A Theory of Contractual Obligation*. Harvard University Press.

Geisst, Charles. 2023. *Just Price in the Markets: A History*. Yale University Press.

Goodin, Robert E. 1991. "Compensation and Redistribution." *NOMOS* 33: 143-177.

Green, Fergus. 2020. "Legal Transitions without Legitimate Expectations." *Journal of Political Philosophy* 28: 397-420.

Greenspan, Alan and Adrian Wooldridge. 2018. *Capitalism in America: a History*. Penguin.

Hayek, F.A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35: 519-530.

Hsieh, Nien-Hê. 2000. "Moral Desert, Fairness and Legitimate Expectations." *Journal of Political Philosophy* 8: 91-114.

Kelly, Jack. 2023. "Goldman Sachs Predicts 300 Million Jobs Will Be Lost Or Degraded By Artificial Intelligence." *Forbes* (31 March 2023). URL: https://www.forbes.com/sites/jackkelly/2023/03/31/goldman-sachs-predicts-300-million-jobs-will-be-lost-or-degraded-by-artificial-intelligence/?sh=7eb44c5782b4.

Knight, Frank H. 2014. *Risk, Uncertainty and Profit*. Martino Publishing.

Kogelmann, Brian. 2022. "We Must Always Pursue Economic Growth." *Utilitas* 34: 478-492.

Kogelmann, Brian. 2024. "Creative Destruction and the Autonomous Life." *Journal of Business Ethics* 197: 659-671.

Kogelmann, Brian. 2025. "Autonomy, Zoning, and Gentrification." *Politics, Philosophy & Economics* (Online First). Doi: https://doi.org/10.1177/1470594X251330912.

Kogelmann, Brian and Jeffrey Carroll. 2024. "The Moral Status of Pecuniary Externalities." *Journal of Business Ethics* 195: 121-132.

Kogelmann, Brian and Jeffrey Carroll. 2025. "The Ethics of AI Automation: The Importance of Treating Like Cases Alike." *AI and Ethics* 5: 3103-3113.

Koyama, Mark and Jared Rubin. 2022. *How the World Became Rich*. Polity Press.

Kuhlmann, Meinard. 2014. "Explaining Financial Markets in Terms of Complex Systems." *Philosophy of Science* 81: 1117-1130.

Liberto, Hallie. 2018. "Promises and the Backward Reach of Uptake." *American Philosophical Quarterly* 55: 15-26.

Mankiw, Gregory N. 2015. *Principle of Microeconomics*. Cengage Learning.

Markovits, Daniel. 2021. "Philosophy of Contract Law." *Stanford Encyclopedia of Philosophy* (23 November 2021), edited by Edward N. Zalta. URL: https://plato.stanford.edu/entries/contract-law/#BarePromDuti.

Melenovsky, C.M. 2021. "Conventionalism and Legitimate Expectations." *Journal of Moral Philosophy* 18: 108-130.

Meyer, Lukas H. and Santiago Truccone-Borgogno. 2022. "Legitimate Expectations: Assessing Policies of Transformation to a Low-Carbon Society." *Environmental Values* 31: 701-720.

Meyer, Lukas H. and Santiago Truccone. 2025. "Legitimate Expectations and Compensation in Changing Circumstances." *Journal of Ethics* (Online first).

Moore, Margaret. 2017. "Legitimate Expectations and Land." *Moral Philosophy and Politics* 4: 229-255.

Nolte, George. 1994. "General Principles of German and European Administrative Law: A Comparison in Historical Perspective." *Modern Law Review* 57: 191-212.

O'Connor, Clare. 2025. "AI replacing jobs? CEOs sound the alarm for white-collar jobs." *Investor's Business Daily* (8-21-2025). URL: https://www.investors.com/news/ai-replacing-jobs-artificial-intelligence-white-collar-workers/.

Page, Scott E. 2011. *Diversity and Complexity*. Princeton University Press.

Procopio, Joe. 2025. “Anthropic’s CEO said all code will be AI-generated in a year.” *Inc.* (20 March 2025). URL: https://www.inc.com/joe-procopio/anthropics-ceo-said-all-code-will-be-ai-generated-in-a-year/91163367.

Quigley, Travis. 2023. “NIMBYism and Legitimate Expectations.” *Journal of Applied Philosophy* 40: 708-724.

Rawls, John. 1971. *A Theory of Justice*. Harvard University Press.

Rodrik, Dani. 2018. *Straight Talk on Trade*. Princeton University Press.

Scanlon, Thomas. 1990. “Promises and Practices.” *Philosophy & Public Affairs* 19: 199-226.

Schwartz, Alan and Robert Scott. 2003. “Contract Theory and the Limits of Contract Law.” *Yale Law Journal* 113: 541-619.

Schumpeter, Joseph. 2008. *Capitalism, Socialism and Democracy*. Harper Perennial.

Sidgwick, Henry. 1981. *The Methods of Ethics*. Hackett Publishing.

Simmons, A. John. 1996. “Associative Political Obligations.” *Ethics* 106: 247-273.

Smith, Adam and Bruce Yandle. 2014. *Bootleggers and Baptists*. Cato Institute.

Traverso, Silvio, Massimiliano Vatiero, and Enrico Zaninotto. 2023. “Robots and Labor Regulation: a Cross-Country/Cross-Industry Analysis.” *Economics of Innovation and New Technology* 32: 977-999.

Vargo, Stephen L., Kaisa Koskela-Huotari, Steve Baron, Bo Edvardsson, Javier Reynoso, and Maria Colurico. 2017. “A Systems Perspective on Markets—Toward a Research Agenda.” *Journal of Business Research* 79: 260-268.

Vatiero, Massimiliano. 2021. *The Theory of Transaction in Institutional Economics*. Routledge.

Williamson, Oliver. 1979. “Transaction-Cost Economics: The Governance of Contractual Relations.” *The Journal of Law & Economics* 22: 233-261.