A California Effect For Artificial Intelligence

By Henry Josephson

Abstract

The California Effect occurs when California's large market, its capacity to successfully regulate, its preference for stringent standards, the inability of the regulatory target to simply move beyond California's jurisdiction, and non-divisibility of the regulatory target combine to mean that companies adhere to California regulation even outside California's borders. In this paper, I look into three ways in which California could regulate artificial intelligence and ask whether each would produce a *de facto* California Effect. I find it likely (~80%) that regulating training data through data privacy would produce a California Effect. I find it unlikely (~20%) that regulation based on the number of floating-point operations needed to train a model would produce a California Effect. Finally, I find it likely (~80%) that risk-based regulation like that proposed by the European Union would produce a California Effect.

Executive summary

Artificial intelligence has the potential to change the world. It remains to be seen whether that change will be for better or for worse. Indeed, because it is still so new, AI has yet to be meaningfully regulated. In this paper, I explore a few ways that California could regulate AI, paying particular attention to the potential for a "California Effect," which occurs when it is easier for companies to offer a California-compliant product everywhere than to produce two different products — one for California, one for everywhere else. This leads to California-compliant products being offered beyond its borders. I focus specifically on this *de facto* California Effect.

First, I explore the theory behind the California Effect as explained by previous scholars. However, because I use the five criteria explained in this section (market size, regulatory capacity, stringent standards, inelastic targets, and non-divisibility) as a metric with which to gauge the likelihood of each intervention producing a California Effect, readers who are not familiar with the California Effect will probably find this section helpful.

The first regulatory target I examine is the data on which large AI models are trained. Ultimately, I find it likely that regulation taking the rough form of *all models trained on California citizens' personally identifiable information must have x safety features* could produce a California Effect. Since the huge datasets that modern AI models need make it infeasible to determine whether a datum's creator is Californian, AI companies cannot avoid training their models on data created by Californians. Indeed, the term 'personally-identifiable information' is so broad that it is impossible for most models not to train on it. Training data's early place in the creation of a model makes it very expensive not to just use the same model everywhere. California is willing to regulate data privacy, and has recently established a government agency to do so. Finally, there is precedent: previous data protection laws have produced California and Brussels Effects.

The second regulatory target I examine is the computing power with which AI models are trained. Ideally, regulations of this kind would adopt (or at the very least allude to) the basic structure of *if a model takes more than n floating point operations to train, it must have x safety features*. High demand for new models (which, in turn, drives demand for training), the inelasticity of the models themselves, and the relative ease with which I believe a regulatory entity could be established all make a California Effect more likely. However, I believe these are outweighed by the apparent lack of public will to regulate large-scale models and the potential for relatively-late forking (i.e. taking a less-trained snapshot of the model), leading me to conclude that regulating compute is unlikely to produce a California Effect.

Finally, I examine whether there would be a California Effect were California to implement regulation similar to the European Union's proposed AI Act. This section's analysis is based on Charlotte Siegmann and Markus Anderljung's paper *The Brussels Effect and Artificial Intelligence*,[1] in which they argue that certain parts of the European Union's proposed AI Act would have a Brussels Effect.

I argue that, if California adopts similar rules, then regulatory diffusion is likely. This is for the same reasons Anderljung and Siegmann argue it is likely to occur for parts of the EU AI Act and because any regulatory diffusion from California would likely amplify regulatory diffusion from the European Union. I also touch on ways in which a Californian implementation of the EU AI Act could differ from an EU implementation.

I conclude with a few thoughts on fruitful topics for further research.

My hope — and ultimate theory of impact — is that this paper will help policymakers make better-informed decisions about future AI regulations. I hope to encourage those who believe in regulating artificial intelligence to give more attention to the State of California. At the very least, I hope that people with a broader reach than I have in the AI Governance space will read and even build off this work. I hope I can raise their awareness of the California Effect and ensure that they recognize the disproportionate impact it can have in the race to keep artificial intelligence safe.

---

[1] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence."

Contents

I - The Theory Behind the California Effect

In this section, I explain the theory underpinning the California Effect, primarily by summarizing and expanding upon Anu Bradford's 2020 *The Brussels Effect*. Readers who are already familiar with the theory may skip to the second part of this document.

However, I recommend that any reader who is unfamiliar with the theory read this section. This is because the analyses in the latter half use the theoretical framework outlined below to assess how likely each regulatory approach is to produce a California Effect. I also summarize the theory below in table 1.

Table 1: Theory Summarized

Factor	Why it makes a California Effect more likely
1. Market size	There is a higher opportunity cost of leaving larger markets, so firms have more incentive to follow a larger market's strict regulation than a smaller markets' regulations. A California Effect requires strict regulation, and their opportunity cost of leaving means that larger markets can impose higher regulatory costs.
2. Regulatory capacity	California must have the capacity to enforce its regulation. If the regulating authority is stretched too thin to investigate, prosecute, and punish violations, then the law does not *de facto* apply to its fullest extent.
3. Preference for stringent standards	For strict regulation to exist in the first place, there must be a preference for strict standards among lawmakers, which usually implies a preference for strict standards among the people.
4. Inelastic targets	'Elasticity' here refers to the ease with which regulation can be avoided by moving its target. If regulation can be evaded by merely moving the regulatory target elsewhere, then a California Effect is less likely to occur.
5. Non-divisibility	If it is cheaper to differentiate production and create both a California-compliant version of the product and a non-California-compliant version than it is to offer the California-compliant version everywhere, then a California Effect is unlikely to occur.

What Is It?

For most, it may at first appear intuitive that a government's rule-making authority only extends as far as its borders — that is, for example, if one were to cross the border from California to Nevada, CA state law would stop applying to that individual and NV state law would instead take its place, specifically with regards to dictating how one may or may not choose to live. For the most part, this intuition is correct. However, for companies operating in both jurisdictions, there exist circumstances in which it is simpler and cheaper to simply follow the stricter state's regulation in both territories. In such circumstances, even though the stricter state's rules may *technically* only apply within its own borders, the rules still end up influencing the behavior of people outside the state, thanks to economic factors described below.

This phenomenon was first described in regards to California by UC Berkeley Business and Political Science Professor David Vogel in his 1995 book *Trading Up*. Vogel coined the phrase "California Effect" to refer to a phenomenon he saw in which states with less-strict regulations end up following the more-demanding regulations from states like California.[2] Regulation enacted by the European Union can also result in this exchange, instead dubbed the "Brussels Effect."

However, the California Effect and its implications for inter- or intra-territory interactions are not as neat as the overarching definition may imply. The California Effect comes in two forms: *de jure* and *de facto*.

In the *de jure* California Effect, regulatory diffusion occurs when other states adopt California's stricter standards for themselves. Vogel provides an illustrative example:

> The California Effect can be seen literally in the history of American automobile emissions standards. The 1970 Clean Air Act Amendments specifically permitted California the option of enacting stricter emissions standards than those required for the rest of the United States, an option which California chose. Consequently its standards remained stricter than any other state. In 1990, Congress brought national emission standards up to California's and once again permitted California to impose stricter standards. It also gave other states the option of choosing either national or California standards. In 1994 twelve eastern states requested that the federal government permit them to adopt California's new standards. These standards in turn are likely to become the basis for the next round of minimum federal requirements. Thus, [. . .] California has helped make American mobile emissions standards steadily stronger.[3]

---

[2] Some "national consumer and environmental regulations exhibit the California Effect: they have moved in the direction of political jurisdictions with stricter regulatory standards." Vogel, *Trading Up*, 259.
[3] Vogel, *Trading Up,* 259.

With the *de jure* California Effect, other states make their own regulations that mirror extant California regulations for reasons that can range from standardization to public pressure to private pressure.

With the *de facto* California Effect, on the other hand, it need not be the case that other states enact regulations that mirror California's for entities within those other states to follow California's stricter regulations. "Due to its large market and preference for stringent consumer and environmental regulation, California is, at times, effectively able to set the regulatory standards for all other states. Businesses willing to export to California must meet its standards, and the benefits from uniform production give these firms an incentive to apply this same (stringent) standard to their entire production."[4] In other words, under certain assumptions, it can be cheaper for a company to simply offer the California-compliant version of their product everywhere than it would be for that company to produce two separate versions of their product (one California-compliant and one not). Thus, non-California consumers can be protected by California law.

Consider, for example, products which warn that they *contain chemicals known to the State of California to cause cancer and birth defects or other reproductive harm*.[5] Even those who live outside California have probably seen such a label. After all, when California enacted the law requiring such warnings, companies who sold their products in California and other states had three choices: stop selling their products in California, label only the products being sold in California, or label all their products.

Of course, since California is both one of the wealthiest and most-populous states, forgoing its market is hardly an option. Its gross state product of $3.6 trillion in 2022 Q1[6] was the largest of any state.[7] Parallel to its economic might is its population, the largest of any state at 39 million people. The opportunity cost of forgoing such a large market is likely quite high.

Especially with economies of scale, it can be cheaper to add one machine which applies a warning sticker to the assembly line than to create another separate assembly line which only produces California-compliant products or to predict what proportion of product will be sold in California and then selectively apply labels to that proportion. Further,

---

[4] Bradford, *The Brussels Effect*, 5.
[5] This is thanks to California's 1986 ballot proposition 65, which "protects the state's drinking water sources from being contaminated with chemicals known to cause cancer, birth defects or other reproductive harm, and requires businesses to inform Californians about exposures to such chemicals."
[6] See the Bureau of Economic Analysis' "Gross Domestic Product by State, 1st Quarter 2022."
[7] Indeed, were California its own country, its economy would only be smaller than the United States, China, Japan, and Germany. (See the World Bank's GDP data here.)

this is likely cheaper than paying steep penalties which violators face.[8] In this way, market forces create a *de facto* California Effect, affording consumers outside California the protection of California law.

Of course, it is not the case that *every* law enacted by the California state legislature automatically becomes national law — consider e.g. varying state responses to recreational marijuana and abortion. Fortunately, there is a way to predict whether and to what extent laws from one jurisdiction will *de facto* apply in another.

The Five Factors

Though it deals with the European Union instead of California, Anu Bradford's 2020 book *The Brussels Effect* is useful for unpacking the topic at hand.[9] As Bradford notes, because previous discussions of the California Effect have not sufficiently explored its causes,[10] she takes a chapter of her book to "outline the precise conditions that allow an upward regulatory convergence to emerge."[11] In that chapter, Bradford argues that "a careful examination of unilateral regulatory authority suggests that there are five elements underlying the Brussels Effect — *market size*, *regulatory capacity*, *stringent standards*, *inelastic targets*, and *non-divisibility*."[12] The stronger each of the elements for a given regulation, the more likely a California/Brussels Effect.

1. Market Size

A larger market correlates with more economic power. This size ought not be judged absolutely, though — relative market size is a far better indicator.[13] Especially since regulatory diffusion occurs with respect to specific regulations, it is also important to think of each jurisdiction's market for whatever *x* is being regulated. Larger markets tend to have wealthier consumers who want to buy *x*, a greater quantity of consumers

---

[8] As Ann Grimaldi, a San Francisco attorney familiar with prop 65 litigation, warns, "It can be very difficult – and it is certainly expensive – for a business to demonstrate that a Prop 65 warning does not apply to its product. The penalties for failure to comply with Prop 65 can be substantial. Fines up to $2500 per day for each violation can be imposed on an entity that has been found to violate Prop 65. Additional penalties of up to $2500 per day per violation can be imposed if a public prosecutor initiates an unfair business practice claim." Grimaldi, "Enforcement And Penalties For Prop 65 Violations - Ouch."
[9] For those interested, a copy is available online for free at https://academic.oup.com/book/36491.
[10] "The theory underpinning the California Effect recognizes the importance of market size and scale economies as a source of a jurisdiction's external regulatory clout. Yet it fails to acknowledge factors such as regulatory capacity and inelasticity as key components of the theory, and overlooks factors other than scale economies that can prevent a company from producing different varieties for different markets. Thus, the discussion of the Brussels Effect provides a more nuanced theory of the conditions under which a single jurisdiction can exert regulatory influence outside its borders." Bradford, *The Brussels Effect*, 5.
[11] Bradford, *The Brussels Effect*, 5.
[12] Bradford, *The Brussels Effect*, 25.
[13] "Market size is relative. The extent of any state's market power depends on the attractiveness of its consumer market compared to the alternative markets available." Bradford, *The Brussels Effect*, 26.

who want to buy it, or both. In other words, not selling x in a larger market has a larger opportunity cost than not selling x in a smaller market. This means that, in theory, whoever sells x should be willing to pay more or tolerate a higher regulatory burden for the right to sell x in the larger market than in the smaller market.

This emphasis on relative market size has important implications. First, it means that, everything else equal, jurisdictions with larger markets for x can impose stricter regulations on x than jurisdictions with smaller markets for x before the firm producing x decides that the costs of compliance if they stay outweigh the opportunity costs if they leave. Second, it means that California or Brussels Effects can only occur when jurisdictions regulate things within their respective markets. That is, these effects are only really relevant when they concern some x which can be denied access to the jurisdiction's large market for x. Third, and perhaps most foundationally, it means that the Brussels and California Effects can produce regulatory diffusion when they are regulating things which are bought and sold. As Bradford explains,

> The EU has little leverage over targets of regulation that are not subject to market access. Consider human rights, an area in which the EU has both regulatory capacity and a strong preference to pursue high levels of protection. However, the EU has not been particularly successful in exporting its human rights norms or democratic values to countries in North Africa or the Middle East, which lie outside of its direct sphere of influence. For example, signing a human rights treaty can be a condition for a trade agreement with the EU. Enforcing the treaty is another matter. It is much easier for the EU to deny market access to a product that does not meet EU standards than it is to police international practices that involve individuals who never enter the European market.[14]

Ultimately, "the EU derives its power [to have some of its regulations *de facto* shape behavior beyond its borders] from its ability to offer conditional access to its large and valuable market. Thus, the jurisdiction's ability to leverage its large market size remains the foundational condition that sustains the Brussels Effect."[15]

2. Regulatory Capacity

If a jurisdiction wants its regulations to shape behavior outside its borders, that jurisdiction must have sufficient regulatory capacity to create and enforce regulations. As Bradford notes, "this requires both regulatory expertise and resources."[16] This is certainly intuitive — after all, inexpertly-designed regulations can end up being so restrictive that the cost of compliance overwhelms the opportunity cost of just leaving the jurisdiction. Further, it makes sense that, for a regulation to be effective, the entity

---

[14] Bradford, *The Brussels Effect*, 30.
[15] Bradford, *The Brussels Effect*, 30.
[16] Bradford, *The Brussels Effect,* 31.

enforcing it needs both the capacity to both detect when firms violate that regulation and the capacity to punish violations with enormous fines and by excluding them from the market.[17] That sanctioning authority is useful not only in that it can force noncompliant firms into compliance, but also because it can serve as a powerful deterrent against noncompliance in the first place. After all, it would make sense that firms are more likely to follow EU regulations on data protection when they would be fined 4% of their global revenue if they do not.[18]

Interestingly, jurisdictions with strong regulatory capacity are often the very same jurisdictions which have a propensity to enact strict regulation, "as jurisdictions that have the political will to adopt stringent regulations also often deploy that same political will to build strong regulatory institutions."[19] Of course, just because a jurisdiction has the political will to build such institutions does not necessarily imply that such a jurisdiction has "the technical expertise or the financial resources to build the requisite regulatory capacity."[20] This is why large but still-developing economies are less likely to be regulatory leaders and why economies like California and the European Union, which have extensive regulatory experience,[21] are more likely to set regulatory trends.

3. Stringent Standards

As noted above, "even significant regulatory capacity by a large market does not guarantee regulatory influence unless such regulatory capacity is supplemented with the political will to deploy it."[22] Though of course some massively unpopular regulations exist, it is almost always the case that some kind of popular will must support a regulation for it to be enacted. This makes sense — after all, it is more difficult to enforce regulation which lacks public support.[23] This public support:

---

[17] "Only those jurisdictions with the ability to inflict significant costs by excluding noncomplying firms from their markets are able to force regulatory adjustments and incentivize compliance." Bradford, *The Brussels Effect*, 31. Sufficiently-high fines can have a comparable effect, since a fine high enough to outweigh the firm's profit in the region would make the firm lose money in that region.

[18] Regulation (EU) 679/16, art. 83, 2016 O.J. (L119) 1.

[19] Bradford, *The Brussels Effect*, 31.

[20] Bradford, *The Brussels Effect*, 31.

[21] Researchers at George Mason University's Mercatus Center have analyzed the regulatory texts of all fifty states and found that California's corpus is both the longest at around twenty-two million words, compared to the national average of around nine million, as well as the most restrictive (quantified by measuring how often certain words like *shall*, *must*, *may not*, et cetera occur). Their data are available at https://www.quantgov.org/state-regdata-definitive-edition, and are in the works cited under *QuantGov*.

    The EU, in turn, "has seen a rise in the role of the regulatory state, as the institutional developments that accompanied the creation of the single market have bestowed the EU with substantial regulatory capacity." Bradford, *The Brussels Effect*, 31. Bradford spends the next five pages exploring how the EU's regulatory capacity has increased over time — interested readers should consult her text.

[22] Bradford, *The Brussels Effect*, 37.

[23] If the people in a jurisdiction are willing to buy something that does not comply with that jurisdiction's regulation, then multinational producers are less incentivized to only produce compliant versions of that product. This switch, of course, is necessary for the regulation to *de facto* apply beyond that jurisdiction.

is more likely to be found in [jurisdictions] with high levels of income.[24] Wealthier countries can better afford pursuing environmental and consumer protection, even at the expense of the profitability of their firms, whereas less-wealthy countries remain more sensitive to the costs of regulation that constrain business activity and hence limit economic growth. This lower tolerance for the costs of stringent rules, together with their lack of regulatory capacity, explains why emerging markets are unlikely to exercise rule-making power that would match their growing market anytime soon.[25]

This dependence on public support introduces an element of variability. After all, public opinion varies with current events. Indeed, Bradford draws a connection between 'triggers' which spark fear of the thing to be regulated and successfully-enacted regulation.[26] Further, public opinion on regulation can be cultural. Here, Bradford emphasizes the difference between European and American regulatory norms, explaining that "The EU does not share the US' reliance on private litigation and tort liability rules to deter firms from placing unsafe or otherwise harmful products on the market."[27]

4. Inelastic Targets

When Bradford says that 'inelastic targets' are prerequisites for regulatory diffusion, she means that it must not be possible to take the regulatory target and move it elsewhere.

An example of an elastic target would be helpful, here. The European Union has enacted strict regulations about hazardous waste disposal.[28] However, because it is relatively easy to put that waste on a boat and ship it to a country with less-restrictive waste disposal regulation, "illegal transfers of hazardous waste remain common as producers have considerable incentive to evade costly regulations by finding jurisdictions that do not enforce waste management standards."[29] Similarly, under maritime law, ships can fly 'flags of convenience,' registering their ship in whichever nation provides the most

---

[24] Here, in her own footnote, Bradford directs readers to Guasch and Hahn, "The Costs and Benefits of Regulation," and Elliott, Regens, and Seldon, "Exploring Variation in Public Support for Environmental Protection." Both are cited below.

[25] Bradford, *The Brussels Effect*, 37.

[26] "While Americans experienced a cascade of alarming news about various such risks from the 1960s until the 1990s—ranging from contaminated cranberries, thalidomide, mercury-contaminated fish, or large oil spills—some widely published disagreements about those alarm bells had eroded their salience, tempering the demand for further regulations [in the US] by the 1990s." Bradford, *The Brussels Effect*, 38.

[27] Bradford, *The Brussels Effect*, 41. Further, "[The US] is accustomed to limiting the government's ability to exert social control in favor of reserving a larger role for private litigants. Private enforcement fits better with the US' individualistic tradition and culture of litigation. The United States also recognizes that government regulation can burden even harmless activity, making the tort system more appealing as it limits the liability to instances where actual harm has occurred." Bradford, *The Brussels Effect*, 42.

[28] O'Niell, "Dynamics of Regulatory Change," 156.

[29] Bradford, *The Brussels Effect*, 30. Funnily enough, this same sentence re-occurs, verbatim, on p52.

favorable regulations — or, more often, the lack thereof.[30] Because one can simply move their trash elsewhere or register their ship in a different country, trash and ships are elastic regulatory targets. Capital, too, is famously elastic — think back to the Panama Papers, which confirmed that the world's wealthy simply move their money to countries with laxer tax laws so they can pay less in taxes.[31]

So, if the regulatory target cannot easily be moved to another, less-restrictive jurisdiction, then it is likely inelastic enough for a California or Brussels Effect to occur. Conversely, if simply moving the target of the regulation outside the regulator's borders is sufficient to avoid regulation, then such a target is too elastic for a California Effect to occur.

5. Non-Divisibility

If the target of the regulation can easily be split into a version for sale in California and a version for sale everywhere else, then regulating that thing is unlikely to produce regulatory diffusion. After all, regulations are only relevant beyond their borders when companies which operate across those borders "decide to apply this new standard to its products or conduct worldwide."[32] This is another reason why relative market size is so important — if a state makes up a greater proportion of the market, companies will probably prioritize complying with that state's regulations over complying with a smaller state's regulations.

Bradford distinguishes a few different types of divisibility: legal, technical, and economic. A classic example of legal non-divisibility is corporate mergers, since corporations cannot merge in one jurisdiction and stay separate in another. Thus, if the strictest jurisdiction rejects a merger, that rejection *de facto* applies everywhere.

Technical non-divisibility is a bit more straightforward. It occurs when, for technical reasons, it costs more to disentangle the production of a good or service across multiple markets than it costs to just comply with the regulation everywhere.

> "For example, to operate in the EU, Google has to amend its data storage and other business practices to conform to European data protection standards. Given the difficulty of determining with certainty whether a particular user is a "European data subject," Google cannot easily isolate its data collection for the EU. As a result, Google adopts a strategy whereby the company adjusts its global operations to the most demanding EU standard."[33]

---

[30] Often Liberia, Panama, or Honduras. For a maritime law tangent, see Sturmy, "Flags of Convenience."
[31] Fitzgibbon and Hudson, "Five Years Later, Panama Papers Still Having a Big Impact - ICIJ."
[32] Bradford, *The Brussels Effect*, 53.
[33] Bradford, *The Brussels Effect*, 57.

Finally, economic non-divisibility occurs when economies of scale or high initial costs make it unwieldy to produce different products for each regulatory standard instead of just making one that fits the highest standard (which includes all the rest). Scale is not the only benefit of consistent production: it also helps branding, it means firms need not predict demand for each market, it protects reputation in the more-strictly regulated market, and can signal consumers in less-regulated markets that the firm cares about their safety.

How Is This AI-Relevant?

There are a variety of reasons one may want to regulate artificial intelligence: to protect civil rights,[34] to ensure data privacy,[35] and to require AIs be aligned with human values,[36] to name a few. Regardless of the intent behind it, though, regulation is generally difficult to pass. Even putting aside the lack of methods to ensure an artificial intelligence's values are aligned with humans', it can still be important to know how to enact such regulation in an iron-clad way as AI alignment research evolves, and more near-term regulation can slow capabilities progress while simultaneously protecting civil rights.

Of course, I do not want to imply that regulation is only ever a good thing. Not only are there the classic arguments on how regulation stifles innovation,[37] but it may be the case that firms can only tolerate so much regulation,[38] and that governments can only regulate the most important things in order to stay under that threshold. Heavier regulation in the US could also cause firms to disperse globally, which undermines the whole point of heavier regulation.

With that said, even stipulating that regulating AI is a good idea does not automatically make it easy to enact. How easy it is to enact regulation is (roughly) inversely proportional to that regulation's scope.

In this way specifically, the California Effect can be a force-multiplier. If regulation is crafted to ensure it yields regulatory diffusion, it means that proponents of a regulation can spend relatively less time and money for more impact than that amount of time and money could otherwise buy. With a California Effect specifically, it means that groups interested in enacting regulation on a large scale can have the benefits of

---

[34] "Urgent Action Needed over Artificial Intelligence Risks to Human Rights."
[35] Pearce, "Beware the Privacy Violations in Artificial Intelligence Applications."
[36] This is a much bigger problem than many realize. To learn more about this threat, I recommend Richard Ngo's "AGI Safety from First Principles." The paper is conversational, compelling, and accessible even to those without machine learning backgrounds.
[37] Explored in Relihan, "Will Regulating Big Tech Stifle Innovation?"
[38] "A business can tolerate only so much regulation before overhead increases exorbitantly…" Barton, "Behind the Legal Explosion."

farther-reaching regulation while only expending the costs associated with otherwise-narrower regulation. With reference to this paper specifically, one would only need to convince the California state legislature — which, though no small task, is easier than convincing all of Congress — to have a national impact.

Especially if someone wants to maximize the impact they can have with each dollar they spend, the California Effect means that they have the chance to get several states for the price of one. The California Effect is a force-multiplier.

Why California?

Even discounting the outsized effects of potential regulatory diffusion,[39] California is uniquely suited to be the first state to enact landmark AI regulation.

California not only enacts the most regulation of any state in the nation[40] — California has often been one of the first states to enact other technology-related legislation. In the last few years alone, California was the third state[41] to ban facial recognition use by law enforcement,[42] the first to pass landmark data privacy regulation,[43] and is poised to be the first state to regulate AI in hiring.[44] In this way, it looks to me as though, if any state would regulate a powerful new technology with the potential to be misused, it would be California.

Indeed, California has shown willingness to regulate artificial intelligence in the past, with — as I note below — some proposed legislation making it as far as the Governor's desk before being vetoed.[45]

---

[39] Which, admittedly, are the main reasons I think California is important with respect to AI regulation.
[40] See note 21 — California has the longest collection of laws of any state in the nation, and its laws have the most restrictions.
[41] After New Hampshire, per https://www.banfacialrecognition.com/map/. Depending on how you count Illinois' Biometric Information Privacy Act, California might actually be the third. Since BIPA requires affirmative consent from the person whose biometric information is being used instead of banning it outright, though, I do not count it as a facial recognition ban.
[42] Metz, "California Lawmakers Ban Facial-Recognition Software from Police Body Cams."
[43] See cited below Stephens, "California Consumer Privacy Act" and Millar and Marshall, "The State of U.S. State Privacy Laws." Though I have anecdotal evidence for a *de facto* California Effect for the CCPA (i.e. that I started to be allowed to opt out of cookies on US websites around when the CCPA went into effect), I could not find a formal analysis of whether the CCPA actually had a California Effect. The extent to which the CCPA (and its follow-up, the CPRA) are producing a *de jure* effect is also unsure — as Millar and Marshall note, though Virginia, Colorado, Utah, and Connecticut have all passed privacy laws similar to California's, they often have more in common with the European GDPR. They also note that: "the CCPA is currently the only one of the five new state laws that allows a private right of action, and the right is limited to breaches of "personal information" (as that term is defined in a separate California data breach notification law, which is more narrowly defined than the term "personal information" in the CCPA)."
[44] Betts and Ochs, "California's Draft Regulations Spotlight Artificial Intelligence Tools' Potential to Lead to Discrimination Claims."
[45] The CITRIS Policy Lab maintains a database of federal and California AI-related legislation.

Further, California's unique position as the United States' technology hub makes California especially relevant here. To me, the fact that such a large proportion of the AI innovation happening in the US is happening in Silicon Valley[46] means that any Californian AI regulation would still be impactful even if regulatory diffusion does not occur. It would send a powerful signal — I once compared it to Iowa regulating new corn technology or West Virginia suddenly limiting coal production.[47]

II - Applying a California Effect to Artificial Intelligence

Below, I consider the potential for a California Effect from regulating the data on which models are trained, regulating models based on how much computing power they use to train, and regulating models based on risk.

As I note later, further research could consider whether it is feasible to require firms to receive some sort of qualification or license before they can do large AI training runs, or whether it would make sense to impose a higher sales tax on transactions associated with less-safe models, regardless of how "less-safe" is ultimately defined.

Would regulating training data cause regulatory diffusion?

In this section, I look into whether regulating the data on which large AI models are trained could produce a California Effect. I find it likely that regulation taking the rough form of *all models trained on California citizens' personally-identifiable information must have x safety features* could produce a California Effect.

I use the California Consumer Privacy Act's definition of PII, which is any "information that identifies, relates to, describes, is capable of being associated with, or could reasonably be linked, directly or indirectly, [to a particular consumer] or household."[48] This is a rather broad definition — with enough time, effort, and the right tools, a worryingly-large amount of data can be traced back to the individual who created it.

I believe that such regulation could produce a California Effect for five main reasons:

---

[46] I do not have a specific metric beyond intuitions about Silicon Valley and the knowledge that Google, Meta, and OpenAI are all headquartered in California, not to mention the countless smaller companies throughout the state and the non-CA-based tech companies with offices in San Francisco and Los Angeles.
[47] MarkusAnderljung and Charlotte, "Supplement to 'The Brussels Effect and AI: How EU AI Regulation Will Impact the Global AI Market.'" As I note in the linked comment, this is a slightly-flawed analogy, "since coal/corn is bigger for WV/IA than AI is for CA."
[48] https://www.oag.ca.gov/privacy/ccpa

First, the huge datasets that modern AI models need make it infeasible to determine whether a datum's creator is Californian, so AI companies cannot avoid training their models on data created by Californians. Second, the term 'personally-identifiable information' is so broad that it is impossible for most models[49] not to train on it. Third, training data's early place in the creation of a model makes it very expensive not to just use the same model everywhere. Fourth, California is willing to regulate data privacy, and has recently established a government agency to do so. Fifth, previous data protection laws have produced California and Brussels Effects.

Table 2: Regulating training data

Factor	Application	Does it make a California Effect more or less likely?
1. Market size	California has significant markets, both generally and for AI training data.	More.
2. Regulatory capacity	The California Privacy Protection Agency already exists and serves to regulate personal data.	More.
3. Preference for stringent standards	The people of California have voted twice in the past five years to enact data privacy regulation.	More.
4. Inelastic targets	The sheer number of Californians and the difficulty of filtering them from training data means that firms cannot avoid training their model on data from Californians, no matter where they do this training. Many AI models (e.g. image recognition, language processing) rely on training data that necessarily includes data that can be personally identifiable.	More.
5. Non-divisibility	Though modern research tools can make it simple to glean personally-identifiable information from individual data, sheer scale and the unlikelihood of anybody	More.

---

[49] That is, this likely is not the case for, say, AlphaGo, which only trains on chess inputs, which are unlikely to be personally-identifiable. However, PII would be unavoidable for, e.g., language models, which need mountains of text containing personal blogs, Tweets, etc., any of which could be used to deduce the identity of the person who created it.

	clearly noting their state citizenship make it prohibitively difficult to remove Californians' PII from a dataset.	

Models use a lot of training data

As background, AI models hoping to do such complex tasks as processing natural language or identifying images need "massive quantities of data" to achieve anywhere near state-of-the art performance, "with the overarching rule-of-thumb being 'the more data the better.' "[50]

Take GPT-3 as an example. GPT-3 trained on Common Crawl, a dataset "constituting nearly a trillion words"[51] scraped from websites over a course of three years, augmented with the entirety of English-language Wikipedia, the content of every link that had ever been posted to Reddit and received more than three upvotes,[52] and two more proprietary datasets, Books1 and Books2. This is internet-scale data.

Other models, like the Beijing Academy of Artificial Intelligence's 2021 Wu Dao, Google's 2022 LaMDA, and Meta's 2022 Atlas all follow that paradigm and also use massive amounts of training data. (Respectively, they used 4.9 terabytes of English and Chinese text and images,[53] 1.56 trillion words compiled from various datasets,[54] and a combination of English Wikipedia with ten years of Common Crawl data.[55]) As the inconsistent units show, there is a significant "lack of transparency in data collection," since the status quo has a dearth of "clear communication of the ingredients and procedures that make up ML projects with the public."[56] Further, recent literature has suggested that these mammoth amounts of data still are not enough to efficiently train models of such scale.[57]

One consequence of the enormous amount of data being collected is that it is quite expensive — often prohibitively so — for humans to go through it all. This means that

---

[50] Bommasani et al., "On the Opportunities and Risks of Foundation Models," 101. Citations removed.
[51] Brown et al., "Language Models Are Few-Shot Learners," 8.
[52] The links from Reddit were collectd in a dataset called WebText, first described in Radford et al., "Language Models Are Unsupervised Multitask Learners," 3, another paper from researchers at OpenAI.
[53] Feng, "Beijing-Funded AI Language Model Tops Google and OpenAI in Raw Numbers."
[54] Thoppilan et al., "LaMDA," 47.
[55] Izacard et al., "Few-Shot Learning with Retrieval Augmented Language Models," 10.
[56] Jo and Gebru, "Lessons from Archives," 312.
[57] Hoffmann et al., "Training Compute-Optimal Large Language Models."

filtering is often based on heuristics, e.g. OpenAI choosing Reddit links with more than three upvotes.[58]

This may not always be the case,[59] though, as companies grow more and more willing to spend increasing time, effort, and money to ensure their models are trained on good data. Even in the limiting case, in which companies hire individuals to manually comb through their enormous datasets — or train AI models to do the same — I still do not find it likely that such efforts could remove all Californians' data from such training sets, for the simple reason that most media posted on the internet does not come with an explicit location tag. There does not seem to be a reliable way to accurately determine whether a piece of PII identifies a Californian or a citizen of another state. Even if it were always possible to identify the location from which something was posted, this does not take into account Californian citizens on vacation in Oregon or in college

Crucially, though, the fact that someone's location is difficult to deduce does not mean that other personally-identifiable information is absent. While certain types of PII fit patterns that can be caught with simple regular expressions and hardcoded heuristics, the very nature of PII is diverse, contextual, and prolific. Especially with a definition as broad as the CCPA's, it would be difficult *not* to train models on data that could be indirectly used to find someone's name. Indeed, much — though admittedly not all — training data "permits the identity of an individual to be directly or indirectly inferred."[60] Modern tools can make it simple to identify the names of who produced which datum.[61] Especially for large language models, which rely on text generated by actual humans, it is not feasible to solely train on data that is not personally identifiable.

Applying the five factors

Market Size

California certainly has a massive absolute market: it had a gross state product of $3.6 trillion in the first quarter of 2022,[62] the largest of any state. Parallel to its economic might is its population, the largest of any state at 39 million people. California is not

---

[58] Radford et al., "Language Models Are Unsupervised Multitask Learners," 3. They, of course, remove duplicate documents from the training data set (after all, people post Wikipedia links to Reddit). The researchers note on the same page that they used the Reddit upvote threshold as a convenient proxy heuristic indicator for whether other users found the link interesting, educational, or just funny" because "**manually filtering a full web scrape would be exceptionally expensive**." (emphasis mine.)

[59] Thanks to Markus Anderljung for raising this point when reading an earlier version of this document.

[60] Department of Homeland Security, "What Is Personally Identifiable Information?"

[61] Prabhu and Birhane, "Large Image Datasets" explores this, and I discuss it further later on this page.

[62] See the Bureau of Economic Analysis' "Gross Domestic Product by State, 1st Quarter 2022."

only the largest state, though — it is also the most diverse,[63] which ensures that, regardless of what exactly is being sold, there likely exist Californians who want to buy it. Further, California is America's technology hub and is home to an enormous portion of the American economy as a whole. Los Angeles and San Francisco are massive hubs for corporations, and the state itself is home to more than 10% of 2021's Fortune 500.[64]

The relative market for training data specifically is weird, since most of it is available for free on the internet.[65] After all, especially when talking about data for training language-based foundation models, the internet *is* the data. Though there certainly exist proprietary datasets, (see, for example, WebText in notes 52 and 58), most datasets are standardized and available online to anyone who wants them.[66] These datasets are huge, and they usually consist of publicly-accessible data scraped from the internet. (Usually, researchers try to anonymize it, but sheer scale means that the effort is rarely enough.[67])

Importantly, this means that most data being used to train these models is made by everyday internet users. For large language models, these data are blog posts, Tweets, book reviews, forum comments, and most other text ever uploaded to the internet. For image classification, image generation, deepfake detection, and other image processing tasks, the industry standard dataset is ImageNet,[68] which contains 14,197,122 labeled images obtained by [querying] several image search engines"[69] and collecting the results. Of course, this method means that "these images are obtained without consent

---

[63] McCann, "Most & Least Diverse States in America." A great example of this which always seems to surprise people is that, despite its reputation as an incredibly blue state, more Californians voted for Trump than did Texans, Floridians, or indeed citizens of *any* other state. (source) This is possible because California is huge.

[64] https://www.statista.com/statistics/303696/us-fortune-500-companies-by-state/

[65] Notably, though, "top results across technical benchmarks have increasingly relied on the use of extra training data [beyond what is publicly available] to set new state-of-the-art results. As of 2021, 9 state-of-the-art AI systems out of the 10 benchmarks in this report are trained with extra data." (AI Index Report 51). This means that there is definitely a market for data beyond what is publicly available, but it is important to note that this is mostly in data collection, not data creation. It also means that the top-of-the-line models are not forgoing public datasets, just supplementing them with proprietary data (i.e. Google training their models on the text of every book in Google Books).

[66] See chapter two of the Stanford Center for Human-Centered AI's 2022 AI Index Report, which lists the datasets used to train and test most state-of-the-art models.

[67] For example, when datasets contain images of people's faces, reverse image search tools often make it trivial to put names to the faces. See Prabhu and Birhane, "Large Image Datasets" for more.

[68] Prabhu and Birhane note that "Although ImageNet was created over a decade ago, it remains one of the most influential and powerful image databases available today. Its power and magnitude is matched by its unprecedented societal impact."

[69] Deng et al., "ImageNet," 251.

or awareness of the individuals or [institutional review board] approval for collection."[70] For video, datasets consist of videos downloaded from Youtube[71] or search results.[72]

As OpenAI noted, filtering a web scrape can be very expensive,[73] particularly when there is no way to make a computer program do it. This will be important in a moment.

Regulatory Capacity

California has the capacity to regulate training data.

In 2020, Californians voted 56.2% to 43.8%[74] to enact Proposition 24, creating the California Privacy Rights Act. The CPRA was voted in because voters did not think California's already-existing data privacy law, the California Consumer Privacy Act, was strict enough. The CCPA, when enacted back in 2018, made California the first US state with data privacy laws.[75] The CCPA, though, was a compromise between the tech companies themselves and a popular movement which wanted even stricter rules. The people voted the CPRA into law because they did not think the CCPA went far enough.[76] The California Privacy Rights Act created the California Privacy Protection Agency to actually implement and enforce California's data privacy laws. The CPPA can issue regulations, update them, and take action against businesses which violate them.[77]

Though enforcement has yet to actually begin (the CPPA's regulations go into effect on January 1, 2023), it looks like the CPPA will have both the expertise and capacity to actually regulate data privacy. The Agency is led by a five-member board which looks

---

[70] Prabhu and Birhane, "Large Image Datasets," 1.

[71] E.g. Kinetics-700, a dataset which the AI Index Report notes on p21 is the state-of-the-art in video activity recognition, "includes 650,000 large-scale, high-quality video clips from YouTube that display a wide range of human activities." Neither the Index Report nor the paper announcing Kinetics-700 make any note of requiring consent from — let alone notifying — the creators of the videos they use.

[72] E.g. ActivityNet, a dataset which the AI Index Report notes on p23 is the state-of-the-art in video temporal action realization, "contains 700 hours of videos of humans doing 200 different activities" (Index Report p23). Similarly to Kinetics-700, this data is collected from web searches. See Heilbron and Niebles, "Collecting and Annotating Human Activities in Web Videos."

[73] The exact quote is "manually filtering a full web scrape would be exceptionally expensive," from Radford et al., "Language Models Are Unsupervised Multitask Learners," 3. Of course, it is less expensive to get a robot to do it. As I noted above, though, PII, by definition, is nuanced and pervasive to such an extent that it cannot easily be rooted out from a dataset by humans, let alone algorithms.

[74] Padilla, "Statement of Vote, General Election, November 3, 2020."

[75] Meyers and Ulloa, "California Lawmakers Agree to New Consumer Privacy Rules."

[76] PrivacyRights.org, "California Privacy Rights Act: An Overview."

[77] See https://cppa.ca.gov/faq.html

quite competent.[78] The board has met frequently to draft regulations,[79] and they released a draft a few months ago.[80]

Though data used for training AI models does not currently fall under the CPPA's purview, it is certainly adjacent. This leads me to believe that, if required to, the CPPA could regulate training data. In the event that the CPPA lacks the expertise to regulate training data, industry experts abound in California in the Computer Science and Law faculties at Stanford, Berkeley, and UCLA.

The infrastructure needed to regulate AI training data is already present, provided such regulation is in the realm of data privacy.

Stringent Standards

Regulating use of personal data for the purpose of ensuring that resulting systems are safe may justifiably be seen as a stretch by voters. However, the popular will to regulate personal data in California certainly exists — or at least, it existed enough two years ago for the people of California to strengthen their existing laws. Of course, it is possible that the 2020 CPRA was enough to satisfy Californians' appetite for privacy protection.

Further, the California State Legislature has shown willingness to regulate Artificial Intelligence, proposing 23 different bills related to AI over the last two legislative sessions.[81] One of these bills, the Artificial Intelligence Bill to Enact the California Artificial Intelligence Act of 2020, which would have created a position in the state's Department of Technology to evaluate the state government's use of AI and potentially regulate it.[82] Before it could take effect, though, Governor Gavin Newsom vetoed it.[83]

---

[78] The board consists of:
1. Jennifer Urban, the board's chair, a professor at Berkeley Law, director of their Technology and Public Policy Clinic, founder of USC's Intellectual Property and Technology Law Clinic. Before joining the CPPA, she taught "interdisciplinary courses in cybersecurity that emphasize how ethical, legal, and economic frameworks enable and constrain security technologies and policies."
2. Chris Thompson, who spent a decade as chief of staff to Senator Dianne Feinstein, and who additionally has lots of non-privacy-related experience working for the CA state government.
3. Angela Sierra, who was previously California's Chief Assistant Attorney General of the Public Rights Division, and as such oversaw CA's prosecution of (and 2019 settlement with) Equifax over their data breach. She has 33 years of experience in the US Department of Justice.
4. Lydia de la Torre, a professor at UCSC Law and ran their Privacy program. She left private practice where she specialized in privacy, data protection, and cybersecurity to join the CPPA.
5. Vinhcent Le, who is a Technology Equity attorney and has a strong history of protecting Californians' data privacy, access to the internet, and protection under the law.

[79] Find minutes of their meetings here.
[80] Priebe et al., "California Privacy Protection Agency Releases Draft of Proposed Regulations."
[81] See them here, along with every AI-related bill that the US Congress has proposed.
[82] Salas, Artificial Intelligence. 2019 AB-594.
[83] Luiz, "Gov. Newsom Vetoes Artificial Intelligence Bill from Assemblyman Rudy Salas."

Other non-data related AI regulations have also been put in place by various levels of the state Government. For example, California's Equal Employment Opportunity Commission has drafted regulations that would limit AI's role in hiring.[84] With that said, this is regulating where AI models can be used — it is not regulating AI models. Though there does exist public will to protect data privacy in California through regulation, and though there does seem to exist public will to regulate artificial intelligence in California, the two do not seem to have combined yet.

Further, Allen Dafoe and Baobao Zhang analyzed American attitudes toward AI in 2019, and found that, when asked to consider "potential policy issues related to AI" survey respondents ranked data privacy highest on "issue importance" and third-highest on "Likelihood of impacting large numbers of people in the U.S. within 10 years."[85] Of course, this survey was not California-specific, but it does show that, in the US as a whole, data privacy is not only an issue — it is an AI-related issue.

Inelastic Targets

When people create things on the internet — remember, the data on which large AI models are trained largely comes from everyday people publishing things on the internet — they generally do not indicate the state in which they live. This means that it can be difficult to discern which fraction of a dataset was created by Californians. As noted above, this means that it is functionally impossible to remove that data from the training set and thus avoid training models data that can be used to identify California citizens.

Further, there are a lot of people from California who are online: only about nine percent of California's forty million people lack regular internet access.[86] I do not know the exact number, but it seems safe to say that a non-negligible fraction of internet content is made by California citizens. By the same logic, a still-larger fraction of English-language content on the internet is made by California citizens.

The specific extent to which Californians are represented on the internet — and thus in these web-crawling datasets — does not really matter for the "inelastic targets" criterion, as long as it is clear that this is a non-negligible portion. The size of that fraction, when combined with the aforementioned difficulty of figuring out which pieces of data are even part of that fraction, makes it functionally impossible for the companies which are handling it to remove every piece of data created by Californians.

[84] Lazzarotti and Yang, "Draft Regulations in California Would Curb Use of AI in Employment."
[85] Zhang and Dafoe, "Artificial Intelligence: American Attitudes and Trends," 18.
[86] Mackovich-Rodriguez, "California Surpasses 90% Internet Access, Low-Income Homes Still Lacking."

Crucially, this is true regardless of whether the model in question is being trained in San Francisco or New York or London — as long as that model is being trained on a sufficiently-large dataset,[87] there are almost certainly some data created by Californians in that dataset. Further, there is no way to reasonably detect or remove that data. This is similar to the Google's worldwide compliance with European data standards:

> "Given the difficulty of determining with certainty whether a particular user is a "European data subject," Google cannot easily isolate its data collection for the EU. As a result, Google adopts a strategy whereby the company adjusts its global operations to the most demanding EU standard."[88]

Indeed, as the above quote makes clear, there is precedent. Data-protection regulations like the European General Data Protection Regulation and aforementioned CCPA/CPPA have already produced Brussels and California Effects, respectively.[89]

As I explained earlier, the key question in this section is whether one could easily remove the thing being regulated from California's jurisdiction. I also explained earlier that it can be cost-prohibitive to go through training data without automation.

Because firms cannot reasonably remove all the Californians in their datasets, they cannot evade regulation based on training data by moving outside California. Data is an inelastic target.

Non-Divisible Production

As Anderljung and Siegmann note in their report about a Brussels Effect for AI,

> "Companies' decisions of whether to offer EU-compliant products outside the EU will largely depend on how fundamental the changes needed to comply with the regulations will be. The more fundamental the changes – the earlier the "fork" in the system – and the costlier it is for the company to maintain two separate products, the more likely they are to choose non-differentiation. In short, early forking often implies high duplication costs which incentivise companies to offer one product globally once they have developed an EU-compliant product."[90]

If companies can cheaply and easily offer two versions of the product — here, a California-compliant version in California and a non-California-compliant version everywhere else — then they probably will.

---

[87] Remember, even 'small' AI models are trained on tens of thousands of pieces of data.
[88] Bradford, *The Brussels Effect*, 57.
[89] See, respectively, asdfasdfasdf and Williams and Irion, "Dream of Californication."
[90] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence," 45.

Artificial intelligence usually has high upfront costs and depends on economies of scale for success,[91] with most of the cost coming in the early stages (paying engineers to build the model itself, collecting and curating training data, and the computing power needed to actually train it). Once the model is complete, it is far cheaper — though not free — to run the model, pay salespeople to sell it, and fine-tune it. This means that, if a regulation forces a company to change something like training, which is early in their AI-creation process, it is probably cheaper for them to just produce one compliant version for use everywhere than to train two different models on two different datasets,[92] one for California and another for everywhere else.

This means that, if training data were regulated, that regulation would probably not induce differentiation.

Conclusion

It seems likely to me that such a regulation of training data (*if the model is trained on data from CA citizens . . .*) would cause a California Effect. After all, California has a relatively-large demand for training data and an even-larger demand for personal and enterprise-level AI models and the efficiency they bring. It has a government agency which seems to have both the expertise and capacity to enforce regulation related to data privacy and security, as well as a population which has voted twice in recent years to enact such data protection. Datasets' monumental size, when combined with the difficulty of discerning whether a single datum is created by a California citizen make it functionally impossible to avoid training a model on data created by Californians. Finally, because collecting training data occurs so early in the AI-creation process, developers would have to fork their projects early to create both a model that is California-compliant and one that is not.

Would regulating training compute cause regulatory diffusion?

Next, I look into whether regulating the computing power with which AI models are trained could produce a California Effect. Ideally, regulations of this kind would adopt (or at the very least allude to) the basic structure of *if a model takes more than n floating point operations to train, it must have x safety features.*

Despite high demand for new models (which, in turn, drives demand for training), despite the inelasticity of the models themselves, and despite the relative ease with which I believe a regulatory entity could be established, the apparent lack of public will

---

[91] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence," 46,

[92] One of which (the one without any data from Californians) would be functionally impossible to create.

to regulate large-scale models and the potential for relatively-late forking (i.e. taking a less-trained snapshot of the model) made it seem unlikely to me that such a regulation would produce a California Effect.

I summarize these results below in table 3.

Table 3: Regulating the training process

Factor	Application	Does it make a California Effect more or less likely?
1. Market size	There exists demand in California for new models, which implies demand for new models' training.	More.
2. Regulatory capacity	Though there does not exist a regulatory agency which would obviously regulate this, it does not seem prohibitively difficult to create one in the future.	More.
3. Preference for stringent standards	There does not seem to be any preference for stricter regulation of AI training.	Less.
4. Inelastic targets	Firms could not avoid the regulation by moving elsewhere, since the regulation would apply based on a characteristic of its training.	More.
5. Non-divisibility	It does not seem prohibitively difficult to train smaller models to evade the regulation, especially if less-capable models were to signal a commitment to safety.	Less.

Models need a lot of training

The specifics of how exactly models are trained are not as important here as understanding the broad requirements. I have mostly been looking into large AI models here, and large models need lots of training. Indeed,

"computer systems are one of the largest bottlenecks to developing [them]. Foundation models are frequently too large to fit in the main memory of a single accelerator (e.g., GPU) and require an immense amount of computation to train (e.g., > [8.64×10^{22} FLOPs] for GPT-3).[93] Additionally, these models will likely get larger over time: for instance, the compute and memory requirements of state-of-the-art language models have grown by three orders of magnitude in the last three years, and are projected to continue growing far faster than hardware capabilities (Figure 19). Once trained, these large models are expensive to perform inference with and difficult to debug, monitor, and maintain in production applications."[94]

Google, Meta, and OpenAI, all have their own supercomputers (OpenAI because their partnership with Microsoft gives them access to Microsoft's Azure supercomputers).[95]

Applying the five factors

Market Size

There is supply-side demand for compute because of how many Californian entities want to meet growing consumer demand for the resulting new AI models. After all, such industry titans as Google,[96] Meta,[97] and OpenAI[98] all call California home. DeepMind may be headquartered in London,[99] but it is still wholly owned by Alphabet. Microsoft is headquartered in Seattle, but has a significant presence in CA, too.[100] The same goes for IBM's research and development,[101] though they are headquartered in New York.[102] That is part of the draw of California's enormous consumer market — the biggest in the US.

Lots of smaller companies are in California, too, though. As a proxy for smaller companies, note that 39 out of Forbes's list of America's 50 Most Promising AI Companies in 2021 are based in California.[103] Each either has their own data center or rents computing time from one of the big players.

---

[93] The original text is "> 1000 petaFLOP/s-days." One petaFLOP/s-day is a day spent doing 10^{15} calculations each second. 1,000 petaFLOP/s-days means doing 10^{15} calculations every second, every day, for a thousand days. Because this is such an abominable unit, I converted it to floating point-operations. For reference, 8.64×10^{22} FLOPs is $10^{15}*3600*24*10^3$.

[94] Bommasani et al., "On the Opportunities and Risks of Foundation Models," 97. Removed internal citation to Brown et al., "Language Models Are Few-Shot Learners."

[95] Shum, Harry. "Advancing Our Ambition to Democratize Artificial Intelligence."

[96] https://en.wikipedia.org/wiki/Googleplex

[97] https://en.wikipedia.org/wiki/Meta_Platforms

[98] https://en.wikipedia.org/wiki/Pioneer_Building_(San_Francisco)

[99] Williams, Greg. "DeepMind Is Opening a Huge New London Headquarters in 2020." On an unrelated note, the article says there's a double-helix staircase, but both photos of it only seem to show one helix.

[100] https://www.microsoft.com/en-us/about/officelocator/all-offices

[101] https://research.ibm.com/labs/

[102] https://en.wikipedia.org/wiki/Armonk,_New_York

[103] Ohnsman, Alan. "AI 50 2021: America's Most Promising Artificial Intelligence Companies."

When I say big players, I mean big players. Amazon Web Services, Microsoft Azure, Google Cloud, and IBM alone combine for 66% of cloud computing market share in 2022.[104]

If there is a company doing cutting-edge AI work, I would be willing to bet a good amount of money that they are doing some of that work in California. Even in the minority of cases where I lose that bet (e.g. with the eleven companies on that Forbes list that are not in California), those companies are almost certainly hoping to sell their product in California. After all, as I noted above, California's size and diversity mean that it has a significant market for essentially everything.

Regulatory Capacity

As of August 2022, California does not currently have any specific organization which would regulate how much computer power AI models use to train. It does not seem too difficult to delegate that authority to a pre-existing department or to create a new one, though.

California does have a Department of Technology,[105] but they are more responsible for the state government's IT support than any actual enforcement.[106] Realistically, if a regulation on the computing power used to train AI models had to be enforced, it would probably fall — like most state-level law enforcement — to some subsidiary of California's Department of Justice. Of course, it would not be difficult to establish a new enforcement agency in the same way that the California Privacy Rights Act created the California Privacy Protection Agency.

Realistically, as technology advances, as models get more efficient, and as the industry moves beyond the current deep learning paradigm, the FLOP threshold at which the regulation applies would have to be updated. For this purpose, it would probably be best to have a panel of experts who meet regularly (say, every other year) to assess whether the threshold ought to change. Luckily, California is home to some of the world's foremost artificial intelligence experts, with a healthy blend of industry (Google, Meta, OpenAI, etc) and academia (Berkeley, Stanford, UCLA, Caltech, etc). I am sure that creating and staffing that hypothetical panel would not be an issue.

---

[104] Kumar, "Cloud Market Share 2022."
[105] https://cdt.ca.gov/
[106] The law establishing the Department of Technology (Government code §11545, available here) lists the duties of the department in section (b). None involve regulating technology that does not belong to the government of California.

Conveniently, papers announcing new AI models seem to always include how much compute the model took to train,[107] so it also does not seem like testing whether new models meet some to-be-determined threshold would be the difficult part — at least for models that actually get published, which is not all of them.[108]

With that said, especially if it were the law, it would not be overly difficult to require firms which make new models to disclose how much compute it took to train that model to some officer of the California state government. Especially if these disclosures are not made publicly available, I could see this holding true for models which have not been released or which companies do not want their competitors to know about — doubly so if there are steep punishments if law enforcement learns that a firm knowingly failed to disclose that information.

I briefly considered regulating data centers (and the large AI models which they are used to train) through a climate lens after coming across a paper about just how energy-intensive data centers are.[109] After all, according to a 2020 project from the Ecole Polytechnique Federale de Lausanne, "one round of training for some of the most complex machine learning models can emit millions of kilograms of carbon dioxide due to the electricity consumed."[110] California has enacted strict climate-related regulations relatively early many times in the past, too.[111] Especially given recent outrage against celebrities' private jet usage,[112] I thought this could be a promising avenue through which to regulate model size. I was wrong.

It turns out, these data centers are probably powered by clean energy. Through purchasing carbon offsets, Google is already carbon-neutral,[113] and they plan to run all their data centers on completely renewable energy by 2030.[114] Microsoft promises to use 100% renewable energy by 2025.[115] Amazon plans the same.[116] It does not feel like there is much the law could reasonably compel them to do with regard to climate change that they are not already doing. Though California may have the regulatory capacity to

---

[107] See, e.g., Brown et al., "Language Models Are Few-Shot Learners," and Chowdhery et al., "PaLM."
[108] It is difficult to estimate what proportion of new AI models are published, since, obviously, I cannot see the models that go unpublished.
[109] Brownlee, Alexander E. I., et al. "Exploring the Accuracy – Energy Trade-off in Machine Learning."
[110] Trébaol, "CUMULATOR — a tool to quantify and report the carbon footprint of machine learning computations," 2.
[111] For example, as I was typing this document, I got a notification from the *New York Times* that "California [is] to ban the sale of new gasoline cars." Harkening back to Vogel's first analysis of the California Effect for auto emissions, the article notes that "more than a dozen other states typically follow California's lead when setting their own auto emissions standards."
[112] See, e.g. Lieber, "As Celebrity Backlash Grows, So Does Overall Private Jet Use"
[113] "Tracking Our Carbon-Free Energy Progress."
[114] "Tracking Our Carbon-Free Energy Progress."
[115] Smith, "Microsoft Will Be Carbon Negative by 2030."
[116] "Sustainability in the Cloud."

regulate large tech companies' data centers' impact on climate change, the companies themselves already seem to be doing so themselves.

In short, California does not have any pre-existing infrastructure with which to enforce a regulation on the amount of computing power used to train AI models. It would be simple for whatever hypothetical bill establishes the regulation to also allocate the funds for a new section of the California Department of Justice to do so, though.

Stringent Standards

Part of the reason I was initially interested in climate-focused regulation is because a key part of the "stringent standards" section is a *preference* for stringent standards. California has definitely demonstrated a preference for strict climate regulation, but a preference for strict regulation of AI models? Not so much.

The Overton window in AI regulation does not include model size. I spent about an hour combing through the internet looking for people discussing it, and I could not find anything serious. Now, the absence of evidence does not imply evidence of absence, but I think it reasonable to expect that, if something has widespread public support, something relevant would show up after twenty minutes of targeted searching. Similarly, the specific methods with which a model is trained, the equipment on which it is trained, and the like all seem to be absent from discussion on AI regulation.

What the discussion does center around, though, are topics like privacy, bias, explainability, and which decisions models are allowed to make. Indeed, actually-implemented AI regulation addresses such concerns, like California's limitation on AI in hiring,[117] ban on facial recognition in police body cameras,[118] and the European Union's AI Act, which places increasingly-strict requirements on increasingly-dangerous uses for AI.[119] Unless popular opinion changes or the issue can be framed in terms of the above (as is the case with training data above), I find it unlikely that there will be popular support for AI regulation.

Inelastic Targets

The requirement that regulation be difficult to flout by simply moving elsewhere is why any compute-based regulation *that depends on where the model is trained* will not have a California Effect.

Any location-based regulation is unlikely to be effective even within California, for the simple reason that it is relatively easy to create and train models outside California's

[117] Lazzarotti and Yang, "Draft Regulations in California Would Curb Use of AI in Employment."
[118] Metz, "California Lawmakers Ban Facial-Recognition Software from Police Body Cams."
[119] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence," 4.

borders. OpenAI would likely be willing to relocate from their San Francisco office if paying to move to a new office in a new state were cheaper than the costs of compliance. Relocation would be even easier for companies like Google, Meta, and the like, which already have myriad offices outside California. When adding the ease and acceptance of remote work relative to before the pandemic, I find that any regulation with a hook similar to *if this model was built in California . . .* would be relatively easy to evade.[120]

The same is even more true when it comes to the actual act of training. Fiber optic cables transmitting at the speed of light mean it is perfectly reasonable to send it to a data center across the border with Nevada. This assumes that all the training for a model is done at a single data center, though. Recent work[121] has shown that it can be not only feasible but far cheaper to train massive models with a decentralized network of GPUs at data centers scattered around the globe. Increasing interconnectedness means that it is trivial to train models outside California.

Funnily enough, that is exactly what Google already does. They do not have any data centers in California,[122] but have two just over the California border in Nevada and another in Oregon. Meta also lacks data centers in California. They have one each in Oregon, Utah, Arizona, and New Mexico, though.[123] One of Amazon's twenty-five worldwide data centers is in Northern California,[124] but it would be trivial to shift any AI work to, say, their Oregon data center should the need arise. The same is true for Microsoft's data centers in Santa Clara and San Jose[125] — their workloads could be distributed (with admittedly-higher latency) to any of the six data centers they have in the continental US alone.

Of course, this problem can be circumvented by making the regulation apply equally to all models, regardless of where they happen to be trained. Though the state of California lacks jurisdiction to regulate interstate commerce, it does have the authority to ban things it finds dangerous or that it believes violate its citizens' rights. The 'inelastic targets' criterion essentially requires that the regulation be impossible to circumvent by merely moving the thing being regulated somewhere else. If the regulation prevents certain-sized models from entering California without relevant safeguards, moving to a different location would not be enough to get around it.

---

[120] Indeed, this is further complicated by the fact that, as an essentially digital item, it can be nigh-impossible to pin down where exactly an artificially-intelligent system was created.
[121] Yuan et al., "Decentralized Training of Foundation Models in Heterogeneous Environments."
[122] Or at least they do not announce any on the website where they list the locations of their data centers.
[123] https://datacenters.fb.com/
[124] "AWS Service Endpoints - AWS General Reference."
[125] "Microsoft 365 Data Locations - Microsoft 365 Enterprise."

Non-Divisible Production

It would be more difficult to restrain a model below a certain threshold of computing power. After all, models have generally gotten larger in search of better performance.[126] It may be true that models with fewer parameters can be trained more compute-efficiently,[127] but that only increases performance for a set budget of computing power. The experts who set the compute threshold would exist to take this into account.

As Anderljung and Siegmann note, the later you can separate the production of the two models, the easier it is for a company to offer different models for different jurisdictions.[128] Not only is this because the expenses in creating an AI model tend to be front-loaded[129] — it is also because earlier forking means that you have to do later parts of the process twice (i.e. once for each model). Late forking is more efficient and less labor-intensive, since it means that adjustments come more in the shape of fine-tuning than in the shape of creating a whole new model.

It is common practice when training a large model to take snapshots every so often in order to gauge how much progress has been made. Of course, these snapshots are necessarily less capable than the final version of the model, but this does not mean that the snapshots are inept. Though it can be prohibitively expensive to train two models, it could be reasonably cheap to train two of them, especially if the firm doing the training can use safety or legal compliance to explain away the dips in performance.

Even so, changes from the current deep learning paradigm, efficiency advances, and better understandings of scaling laws[130] can make it possible to train more-efficient models with less compute than possible today. Such advances can make any compute threshold less effective — unless, of course, the aforementioned panel of experts exists to raise or lower that threshold in response to changes in the industry.

Conclusion

Since there is certainly a large-enough market size, and since it probably would not be too difficult to establish the capacity to enforce such regulations, whether or not

---

[126] Sevilla et al., "Compute Trends Across Three Eras of Machine Learning." Readers may be especially interested in the interactive version of Sevilla et al.'s findings, current as of January 24, 2022, here.
[127] Hoffmann et al., "Training Compute-Optimal Large Language Models."
[128] See p19 above, where I previously referenced this quote.
[129] This makes sense — it generally takes much more compute to train a model than to run a model once it has been trained, less engineering is needed to maintain, occasionally update, and troubleshoot a model that already exists than to create one from scratch, etc.
[130] Including those like Hoffmann et al., "Training Compute-Optimal Large Language Models," which revealed that much-higher efficiency is possible on the same compute budget if model-builders use more training data.

regulating the amount of computing power needed to train a model produces regulatory diffusion would depend on a few things. First, it depends on the extent to which the public prefers stringent standards on model size. In the status quo, that demand does not seem to exist. Second, it depends on whether the regulation can be followed with early forking. Because late forking seems possible (and because such regulation may incentivize AI creators to make late forking work), I do not believe that regulating compute would cause a California Effect.

Would regulating models themselves cause regulatory diffusion?

There are many ways to regulate models themselves

As the heading for this section suggests, there are many avenues through which to regulate models themselves. Especially given that regulating models based on how risky what they are being used for feels reasonable, I decided to focus on whether there would be a California Effect were California to implement regulation similar to the European Union's proposed AI Act. This analysis is based heavily on Anderljung and Siegmann's paper *The Brussels Effect and Artificial Intelligence*,[131] in which they argue that certain parts of the European Union's proposed AI Act would have a Brussels Effect.

I argue that, if California adopts similar regulations, then regulatory diffusion would likely occur. This is for the same reasons Anderljung and Siegmann argue it is likely to occur for parts of the EU AI Act and because any regulatory diffusion from California would likely amplify regulatory diffusion from the European Union.

How does the European regulation work?

The European approach to artificial intelligence regulation is based on risk,

> "classifying AI systems as creating unacceptable, high, limited, or minimal risk. The level of risk is judged by the likelihood that the system may harm specific individuals, potentially violating their fundamental rights. The requirements imposed on systems are related to the level of risk, ranging from prohibitions to the voluntary adoption of codes of conduct. The AIA proposes prohibitions on AI applications that pose "unacceptable risks", including "real-time" remote biometric identification systems used by governments. It requires conformity assessments for "high-risk" AI systems, such as some AI systems deployed in worker management, critical infrastructure operation, border control, remote biometric identification, medical devices, machinery, and other areas. Certain limited-risk AI systems need to comply with transparency rules, requiring that

---

[131] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence."

users are made aware e.g. if they are engaging with AI-generated content that may appear authentic such as chatbots or deepfakes. All other AI systems, termed "minimal risk", face no additional obligations, though providers are encouraged to follow voluntary codes of conduct."[132]

As noted above, the EU AI Act is not a monolith, and it regulates different AI systems in different ways. As such, Anderljung and Siegmann found that some sections were likely to produce a de facto Brussels Effect, while others were not. Since I do not want to spend too much space rehashing their arguments,[133] I will briefly summarize their findings for the European AI Act.

Some parts of the Act, like user-facing transparency requirements for lower-risk AI applications, are unlikely to produce regulatory diffusion, because it is too easy to meet those requirements with a late fork. E.g., if firms are required to disclose that their customer service chatbots are not humans, it is trivial to insert a single line on the website that says "Hi! I'm a customer service bot!"

Anderljung and Siegmann use analogous Californian technology regulation as a case study.[134] As the name suggests, California's 2018 Bot Disclosure Act[135] requires firms using bots to help sell things or influence votes to disclose that they are bots instead of people. Despite two efforts from California senator Dianne Feinstein in 2018[136] and 2019[137] which both died in committee, an analogous law has never been enacted federally. Anecdotally, many bots with which I've interacted have disclosed that they are bots, it is important to keep in mind that any proliferation here could be adequately explained without invoking a California Effect — after all, industry best practices already advocate telling customers that they are speaking to a bot instead of a real person.[138]

Anderljung and Siegmann do expect to see *de facto* regulatory diffusion for AI in products whose production is already standardized throughout the European Union (such as "medical devices, toys, and machinery"[139]), for AI used in "hiring, firing, and worker management," "some general AI systems or foundation models across a wide

---

[132] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence," 12-13.
[133] Please read their paper.
[134] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence," 50.
[135] "Bill Text - SB-1001 Bots: Disclosure."
[136] Fienstein, "S.3127 - Bot Disclosure and Accountability Act of 2018."
[137] "S.2125 - Bot Disclosure and Accountability Act of 2019."
[138] See e.g. ChatBot, "Your Ultimate Chatbot Best Practices Guide," which tells readers to "be transparent by telling users they are talking with a chatbot."
   It is also possible that this is confirmation bias — after all, I am more likely to notice that I am interacting with a bot if it admits to being a bot.
[139] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence," 23.

range of uses and industries," and, "less confidently, the use of AI in the legal sector and the use of biometric identification and categorisation of natural persons."[140]

Anderljung and Siegmann do not believe that other high-risk uses for AI will produce a Brussels Effect, since they are "already regionalized."[141] Controlling traffic, college admissions, grading schoolwork, law enforcement, border security, and other government tasks are already specific to each institution. For example, there is no multinational firm to which universities outsource their admissions decisions — each university tends to have its own admissions department — so, even though regulations affecting AI in university admissions may impact every European university, those European regulations seem unlikely to have any impact on universities outside Europe.

Interestingly enough, something similar is already occurring in California, albeit not with artificial intelligence. The state's flagship University of California system is going completely test-blind in 2021,[142] a decision which, to date, has not been echoed by any other public universities.[143]

Similar justifications hold true for AI use in critical infrastructure (most electric grids are already regionalized, one country's dams generally do not consult another's when deciding how much water to release, etc.) and for AI in the financial sector.

What does this mean for California?

First and foremost, it means that, first of all, if CA enacts sufficiently-similar AI regulation, and if that regulation is enacted after the AI act is (which, given government timelines, seems likely), it'll probably reinforce whatever regulatory diffusion the EU AI act has produced. Such *de facto* regulatory diffusion from Europe's General Data Protection Regulation is already occurring,[144] and it makes sense that *de jure* regulatory diffusion — especially *de jure* diffusion to a jurisdiction whose regulations often produce regulatory diffusion themselves — would beget even more regulatory diffusion in turn. If both California and the European Union were to enact similar regulations, then the diffusion of those regulations could be propelled by a dual California-Brussels Effect.

---

[140] All three quotes from Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence," 50.
[141] Anderljung and Siegmann, "The Brussels Effect and Artificial Intelligence," 50.
[142] del Rio, "University of California Will No Longer Consider SAT and ACT Scores."
[143] The City University of New York has test-blind admissions with regard to the SAT and ACT through 2023, but they still consider New York State's proprietary standardized tests, the Regents exams. See "Testing FAQs."
[144] If you've ever been asked to give affirmative consent to cookies on a website, that's the GDPR, not California's CPRA, which only requires that you be given the opportunity to opt-out of tracking. See Singh, "What You Need to Know about the CCPA and the European Union's GDPR," which notes this difference.

Further, if California were to adopt regulation sufficiently similar to the European AI Act, it would be a powerful signal that jurisdictions on the cutting edge of regulation think such classifications are important. Such a signal would likely also cause such regulation to be associated with safety and good business practices worldwide, hastening regulatory diffusion.

Of course, this assumes that Californian regulation is sufficiently similar to the European AI Act, which has still yet to be finalized.[145] Though there are many different ways in which California may see fit to adapt the EU's Act, one I think would be particularly likely would be the enforcement mechanism. As Bradford notes in *The Brussels Effect*, "The EU does not share the US' reliance on private litigation and tort liability rules to deter firms from placing unsafe or otherwise harmful products on the market."[146] Indeed, though both California and EU privacy law include governmental authority to fine violators, California law is unique in that it also "creates a private right of action and an entitlement to statutory damages for non-compliance."[147]

If anything, creating a private right of action while also creating a government agency helps regulatory diffusion, if only because having more parties interested in suing violators makes it more likely that violators are punished. This is important especially when considering differences between European and American law. In Europe, the European Commission has the authority to initiate investigations, decide them, and then impose punishment.[148] This is not the case in California, where alleged privacy violators are prosecuted by representatives of the people (i.e. the California Department of Justice) in state courts. This means that the European Union is on balance more likely to impose punishment than California is. Creating a private right of action helps even out this imbalance.

Another potential difference between a Californian AI Act and the European AI Act could be in the way the two regulations treat general-purpose AI systems.[149] Though the EU AI Act has yet to be finalized, "the most controversial amendments [have involved how the regulation deals] with general-purpose AI."[150] These 'foundation models' are fine-tuned into a variety of industry applications, and present a potential path to artificial general intelligence.[151] Especially in the case that the final version of the EU law does not regulate general-purpose technologies, regulating general-purpose AI could be

---

[145] See Bertuzzi, "AI Regulation Filled with Thousands of Amendments in the European Parliament," for some things EU lawmakers are discussing.

[146] Bradford, *The Brussels Effect*, 30.

[147] Singh, "What You Need to Know about the CCPA and the European Union's GDPR,"

[148] See Karayanidi, "Does the European Commission Have Too Much Power Enforcing European Competition Law?"

[149] Thank you to Anthony Barrett for reminding me of this.

[150] Bertuzzi, "Leading MEPs Seek Broad Consensus on Regulatory Sandboxes in AI Act."

[151] Fei et al., "Towards Artificial General Intelligence via a Multimodal Foundation Model."

especially important, both because it would have downstream effects for each of the general-purpose model's specific applications, but also because it would regulate a potential precursor to AGI.

Conclusion

Anderljung and Siegmann found it likely that certain parts of the European Union's proposed AI Act would produce a Brussels Effect. For the same reasons that those parts of the regulation are likely to diffuse from Europe, I believe that they are also likely to diffuse from California. Further, should California enact regulation that is sufficiently similar to the EU AI Act, the California and Brussels Effects would likely amplify each other, leading to still more regulatory diffusion. Though a California version of the law could differ from the European version (e.g. by allowing individuals to sue), this would not weaken the potential for diffusion — and could, in fact, strengthen it.

Opportunities for Further Research

There are plenty of opportunities for further research here. Before I list my own, though, I will direct readers to the [list of further research opportunities](#) that Anderljung and Siegmann collected in an announcement for their report on the potential Brussels Effect of the EU AI Act. I am personally choosing to highlight their fourth and sixth bullet points, which I think would be especially effective (the latter especially):

- "Empirical work tracking the extent to which there is likely to be a Brussels Effect. Most of the research on regulatory diffusion focuses on cases where diffusion has already happened. It seems interesting to instead look for leading indicators of regulatory diffusion. For example, you could analyze relevant parliamentary records or conduct interviews, to gain insight into the potential global influence of the EU AI Act, the EU, and legal terms and framings of AI regulation first introduced in the EU discussion leading up to the EU AI Act.
- Work on what good AI regulation looks like from a TAI/AGI perspective seems particularly valuable. Questions include: What systems should be regulated? Should general-purpose systems be a target of regulation? Should regulatory burdens scale with the amount of compute used to train a system? What requirements should be imposed on high-risk systems? Are there AI systems that should be given fiduciary duties?"

Interested readers should also peruse the Centre for Governance of AI's [research agenda](#), which is far more exhaustive than I could ever hope to be.

With other people's suggestions out of the way, I think there is a dearth of research into the impact state governments can have, in artificial intelligence governance but

especially in other cause areas. State and local governments account for a bit less than half of all government spending in the US[152] yet can be far more accessible. Especially in the context of AI governance, I would love to see more research into state-level interventions, anywhere from research funding/grants to tax breaks.

Interestingly, the California Privacy Rights Act gives the California Privacy Protection agency the right to "Issu[e] regulations governing access and opt-out rights with respect to businesses' use of automated decision-making technology, including profiling and requiring businesses' response to access requests to include meaningful information about the logic involved in such decision-making processes, as well as a description of the likely outcome of the process with respect to the consumer."[153] However, their proposed regulations[154] do not seem to mention AI or automated decision-making. Though the CPPA is no longer accepting comments on their proposed regulations, it could be useful to look into what it would take to get them to include AI to a greater extent.

In this same vein, it could be useful to look at instances in the past when other states' regulatory authorities have attempted to regulate online commerce. Were they successful? What would these previous attempts at regulation mean for future attempts to regulate AI? Though I touched upon the California Privacy Protection Agency, it may be the case that such an agency is not the right entity to create and enforce these regulations. Which other agencies, e.g. consumer protection, could effectively regulate AI? This could also be worth looking into at the federal level, too.

It could also be a good idea to require registration for training runs, data collection, or even the entirety of model creation. As such, research into prior attempts to require licenses for the creation and use of new technology (e.g. transportation, research technologies, weapons, etc) could be useful.

---

[152] Data for fiscal year 2019 at "The Federal Budget in 2019: An Infographic." See also "State and Local Expenditures."

[153] Fath and Oberly, "California Privacy Protection Agency Continues Rulemaking Focus on Automated Decision-Making and Profiling in Stakeholder Sessions."

[154] https://cppa.ca.gov/regulations/pdf/20220708_text_proposed_regs.pdf

Works Cited

"About Us - California Privacy Protection Agency (CPPA)." Accessed August 11, 2022. https://cppa.ca.gov/about_us/.

Anderljung, Markus, and Charlotte Siegmann. "The Brussels Effect and Artificial Intelligence," August 16, 2022. https://www.governance.ai/research-paper/brussels-effect-ai.

"Armonk, New York." In *Wikipedia*, August 28, 2022. https://en.wikipedia.org/wiki/Armonk,_New_York.

"AWS Service Endpoints - AWS General Reference." Accessed August 28, 2022. https://docs.aws.amazon.com/general/latest/gr/rande.html.

Barton, John H. "Behind the Legal Explosion." *Stanford Law Review* 27, no. 3 (February 1975): 567. https://doi.org/10.2307/1228327.

Bertuzzi, Luca. "AI Regulation Filled with Thousands of Amendments in the European Parliament." www.euractiv.com, June 2, 2022. https://www.euractiv.com/section/digital/news/ai-regulation-filled-with-thousands-of-amendments-in-the-european-parliament/.

———. "Leading MEPs Seek Broad Consensus on Regulatory Sandboxes in AI Act." www.euractiv.com, September 2, 2022. https://www.euractiv.com/section/digital/news/leading-meps-seek-broad-consensus-on-regulatory-sandboxes-in-ai-act/.

"Bill Text - SB-1001 Bots: Disclosure." Accessed August 28, 2022. https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB1001.

Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, et al. "On the Opportunities and Risks of Foundation Models," 2021. https://doi.org/10.48550/ARXIV.2108.07258.

Bradford, Anu. *The Brussels Effect: How the European Union Rules the World*. 1st ed. Oxford University Press, 2020. https://doi.org/10.1093/oso/9780190088583.001.0001.

Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. "Language Models Are Few-Shot Learners," 2020. https://doi.org/10.48550/ARXIV.2005.14165.

Brownlee, Alexander E.I, Jason Adair, Saemundur O. Haraldsson, and John Jabbo. "Exploring the Accuracy – Energy Trade-off in Machine Learning." In *2021*

IEEE/ACM International Workshop on Genetic Improvement (GI), 11–18. Madrid, Spain: IEEE, 2021. https://doi.org/10.1109/GI52543.2021.00011.

"California Privacy Protection Agency (CPPA)." Accessed August 11, 2022. https://cppa.ca.gov/.

PrivacyRights.org. "California Privacy Rights Act: An Overview," December 10, 2020. https://privacyrights.org/resources/california-privacy-rights-act-overview.

Carreira, Joao, Eric Noland, Chloe Hillier, and Andrew Zisserman. "A Short Note on the Kinetics-700 Human Action Dataset." arXiv, July 15, 2019. http://arxiv.org/abs/1907.06987.

"CDT | CA Dept of Technology." Accessed August 28, 2022. https://cdt.ca.gov/.

Chowdhery, Aakanksha, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, et al. "PaLM: Scaling Language Modeling with Pathways," 2022. https://doi.org/10.48550/ARXIV.2204.02311.

DaFoe, Allan. "AI Governance: A Research Agenda." Center for the Governance of AI, August 27, 2018. https://www.governance.ai/research-paper/agenda.

Davenport, Coral, Lisa Friedman, and Brad Plumer. "California to Ban the Sale of New Gasoline Cars." *The New York Times*, August 24, 2022, sec. Climate. https://www.nytimes.com/2022/08/24/climate/california-gas-cars-emissions.html.

Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. "ImageNet: A Large-Scale Hierarchical Image Database." In 2009 IEEE Conference on Computer Vision and Pattern Recognition, 248–55. Miami, FL: IEEE, 2009. https://doi.org/10.1109/CVPR.2009.5206848.

Department of Homeland Security. "What Is Personally Identifiable Information?," December 18, 2021. https://www.dhs.gov/privacy-training/what-personally-identifiable-information.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding," 2018. https://doi.org/10.48550/ARXIV.1810.04805.

Elliott, Euel, James L. Regens, and Barry J. Seldon. "Exploring Variation in Public Support for Environmental Protection." *Social Science Quarterly* 76, no. 1 (1995): 41–52. http://www.jstor.org/stable/44072588.

Google Data Centers. "Discover Our Data Center Locations." Accessed August 28, 2022. https://www.google.com/about/datacenters/locations/.

Heilbron, Fabian Caba, and Juan Carlos Niebles. "Collecting and Annotating Human Activities in Web Videos." In Proceedings of International Conference on Multimedia Retrieval, 377–84. Glasgow United Kingdom: ACM, 2014. https://doi.org/10.1145/2578726.2578775.

Fath, Kyle R., and David J. Oberly. "California Privacy Protection Agency Continues Rulemaking Focus on Automated Decision-Making and Profiling in Stakeholder Sessions." The National Law Review, May 10, 2022. https://www.natlawreview.com/article/california-privacy-protection-agency-continues-rulemaking-focus-automated-decision.

"Federal & California AI Legislation." CITRIS Policy Lab and the Banatao Institute, February 17, 2022. https://citrispolicylab.org/ailegislation/.

Fei, Nanyi, Zhiwu Lu, Yizhao Gao, Guoxing Yang, Yuqi Huo, Jingyuan Wen, Haoyu Lu, et al. "Towards Artificial General Intelligence via a Multimodal Foundation Model." Nature Communications 13, no. 1 (June 2, 2022): 3094. https://doi.org/10.1038/s41467-022-30761-2.

Feng, Coco. "Beijing-Funded AI Language Model Tops Google and OpenAI in Raw Numbers." South China Morning Post, June 2, 2021. https://www.scmp.com/tech/tech-war/article/3135764/us-china-tech-war-beijing-funded-ai-researchers-surpass-google-and.

Fienstein, Dianne. "S.2125 - Bot Disclosure and Accountability Act of 2019," July 16, 2019. https://www.congress.gov/bill/116th-congress/senate-bill/2125.

———. "S.3127 - Bot Disclosure and Accountability Act of 2018," June 25, 2018. https://www.congress.gov/bill/115th-congress/senate-bill/3127.

Fight for the Future. "See Where Dangerous Facial Recognition Is Being Used, and Learn What You Can Do about It." Ban Facial Recognition. Accessed August 28, 2022. https://www.banfacialrecognition.com/map/.

Fitzgibbon, Will, and Michael Hudson. "Five Years Later, Panama Papers Still Having a Big Impact - ICIJ," April 3, 2021. https://www.icij.org/investigations/panama-papers/five-years-later-panama-papers-still-having-a-big-impact/.

"Frequently Asked Questions (FAQs) - California Privacy Protection Agency (CPPA)." Accessed August 11, 2022. https://cppa.ca.gov/faq.html.

"Full- and Part-Time Legislatures." National Conference of State Legislatures, July 28, 2021. https://www.ncsl.org/research/about-state-legislatures/full-and-part-time-legislatures.aspx.

"Googleplex." In *Wikipedia*, July 24, 2022. https://en.wikipedia.org/wiki/googleplex.

Grimaldi, Ann. "Enforcement And Penalties For Prop 65 Violations - Ouch," August 3, 2018. https://grimaldilawoffices.com/enforcement-and-penalties-for-prop-65-violations-ouch/.

Guasch, J. L., and R. W. Hahn. "The Costs and Benefits of Regulation: Implications for Developing Countries." *The World Bank Research Observer* 14, no. 1 (February 1, 1999): 137–58. https://doi.org/10.1093/wbro/14.1.137.

Haataja, Meeri, and Joanna J. Bryson. "What Costs Should We Expect from the EU's AI Act?" Preprint. SocArXiv, August 27, 2021. https://doi.org/10.31235/osf.io/8nzb4.

Hoffmann, Jordan, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, et al. "Training Compute-Optimal Large Language Models," 2022. https://doi.org/10.48550/ARXIV.2203.15556.

Izacard, Gautier, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. "Few-Shot Learning with Retrieval Augmented Language Models," 2022. https://doi.org/10.48550/ARXIV.2208.03299.

Berkeley Law. "Jennifer Urban." Accessed August 28, 2022. https://www.law.berkeley.edu/our-faculty/faculty-profiles/jennifer-urban/.

Jo, Eun Seo, and Timnit Gebru. "Lessons from Archives: Strategies for Collecting Sociocultural Data in Machine Learning." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 306–16. Barcelona Spain: ACM, 2020. https://doi.org/10.1145/3351095.3372829.

Johnson, Steven, and Nikita Iziev. "A.I. Is Mastering Language. Should We Trust What It Says?" *The New York Times*, April 15, 2022, sec. Magazine. https://www.nytimes.com/2022/04/15/magazine/ai-language.html.

Karayanidi, Milana S. "Does the European Commission Have Too Much Power Enforcing European Competition Law?" *German Law Journal* 12, no. 7 (July 1, 2011): 1446–59. https://doi.org/10.1017/S2071832200017387.

Kumar, Rahul. "Cloud Market Share 2022: An Overview Of Growing Ecosphere,"
     April 22, 2022. https://www.wpoven.com/blog/cloud-market-share/.

IBM Research. "Labs and Locations," June 5, 2018. https://research.ibm.com/labs.

Lazzarotti, Joseph J., and Robert Yang. "Draft Regulations in California Would Curb
     Use of AI, Automated Decision Systems in Employment." The National Law
     Review, April 11, 2022.
     https://www.natlawreview.com/article/draft-regulations-california-would-curb
     -use-ai-automated-decision-systems-employment.

"Legality of Cannabis by U.S. Jurisdiction." In *Wikipedia*, August 27, 2022.
     https://en.wikipedia.org/wiki/Legality_of_cannabis_by_U.S._jurisdiction.

Lieber, Chavie. "As Celebrity Backlash Grows, So Does Overall Private Jet Use."
     Wall Street Journal, August 14, 2022, sec. Style News.
     https://www.wsj.com/articles/private-jet-celebrities-taylor-swift-kylie-jenner-
     backlash-11660345684.

Meta Data Centers. "Locations." Accessed August 28, 2022.
     https://datacenters.fb.com/#locations.

Luiz, Joseph. "Gov. Newsom Vetoes Artificial Intelligence Bill from Assemblyman
     Rudy Salas." *KGET 17*, October 2, 2019.
     https://www.kget.com/news/gov-newsom-vetoes-artificial-intelligence-bill-fro
     m-assemblyman-rudy-salas/.

MarkusAnderljung and Charlotte. "Supplement to 'The Brussels Effect and AI: How
     EU AI Regulation Will Impact the Global AI Market.'" *Effective Altruism
     Forum*, August 16, 2022.
     https://forum.effectivealtruism.org/posts/gJGMFdGqFhs3mKo2s/supplement
     -to-the-brussels-effect-and-ai-how-eu-ai?commentId=cbLZCqZZxaSHbcbSo.

McCann, Adam. "Most & Least Diverse States in America." WalletHub, September
     21, 2021.
     https://wallethub.com/edu/most-least-diverse-states-in-america/38262.

"Meetings - California Privacy Protection Agency (CPPA)." Accessed August 11,
     2022. https://cppa.ca.gov/meetings/.

"Meta Platforms." In *Wikipedia*, August 27, 2022.
     https://en.wikipedia.org/wiki/Meta_Platforms.

Metz, Rachel. "California Lawmakers Ban Facial-Recognition Software from Police
     Body Cams." CNN, September 13, 2019.

https://www.cnn.com/2019/09/12/tech/california-body-cam-facial-recognitio
n-ban/index.html.

"Microsoft 365 Data Locations - Microsoft 365 Enterprise," August 19, 2022.
https://docs.microsoft.com/en-us/microsoft-365/enterprise/o365-data-locatio
ns.

"Microsoft Office Locations List." Microsoft. Accessed August 12, 2022.
https://www.microsoft.com/en-us/about/officelocator/all-offices.

Millar, Sheila A., and Tracy P. Marshall. "The State of U.S. State Privacy Laws: A
Comparison." The National Law Review, May 24, 2022.
https://www.natlawreview.com/article/state-us-state-privacy-laws-comparison
.

Epoch. "ML Input Trends Visualization." Accessed August 28, 2022.
https://epochai.org/mlinputs/visualization.

Myers, John, and Jazmine Ulloa. "California Lawmakers Agree to New Consumer
Privacy Rules That Would Avert Showdown on the November Ballot." Los
Angeles Times, June 22, 2018.
https://www.latimes.com/politics/la-pol-ca-privacy-initiative-legislature-agree
ment-20180621-story.html.

Ngo, Richard. "AGI Safety from First Principles," September 2020.
https://drive.google.com/file/d/1uK7NhdSKprQKZnRjU58X7NLA1auXlWHt/e
dit?usp=embed_facebook.

Ochs, Danielle, and Jennifer Betts. "California's Draft Regulations Spotlight
Artificial Intelligence Tools' Potential to Lead to Discrimination Claims." The
National Law Review, May 13, 2022.
https://www.natlawreview.com/article/california-s-draft-regulations-spotlight-
artificial-intelligence-tools-potential-to.

Ohnsman, Alan. "AI 50 2021: America's Most Promising Artificial Intelligence
Companies." Forbes, April 26, 2021.
https://www.forbes.com/sites/alanohnsman/2021/04/26/ai-50-americas-mos
t-promising-artificial-intelligence-companies/.

O'Neill, Kate. "The Changing Nature of Global Waste Management for the 21st
Century: A Mixed Blessing?" Global Environmental Politics 1, no. 1 (February 1,
2001): 77–98. https://doi.org/10.1162/152638001570642.

Google Sustainability. "Our Sustainability Efforts & Progress." Accessed August 20,
2022. https://sustainability.google/progress/.

Padilla, Alex. "Statement of Vote: General Election," November 3, 2020.
    https://elections.cdn.sos.ca.gov/sov/2020-general/sov/complete-sov.pdf.

Pearce, Guy. "Beware the Privacy Violations in Artificial Intelligence Applications."
    ISACA. Accessed August 28, 2022.
    https://www.isaca.org/resources/news-and-trends/isaca-now-blog/2021/bewa
    re-the-privacy-violations-in-artificial-intelligence-applications.

"Pioneer Building (San Francisco)." In *Wikipedia*, April 11, 2021.
    https://en.wikipedia.org/wiki/Pioneer_Building_(San_Francisco).

Prabhu, Vinay Uday, and Abeba Birhane. "Large Image Datasets: A Pyrrhic Win for
    Computer Vision?" arXiv, July 23, 2020. http://arxiv.org/abs/2006.16923.

"Presidential Election Results: Live Map of 2020 Electoral Votes," August 1, 2022.
    https://www.nbcnews.com/politics/2020-elections/president-results.

Priebe, Jason, Tom Tomaszewski, and Danny Riley. "California Privacy Protection
    Agency Releases Draft of Proposed Regulations to the CPRA." Carpe Datum
    Law, June 27, 2022.
    https://www.carpedatumlaw.com/2022/06/california-privacy-protection-agen
    cy-releases-draft-of-proposed-regulations-to-the-cpra/.

"Proposition 65 FAQs." California Office of Environmental Health Hazard
    Assessment, February 1, 2014.
    https://oehha.ca.gov/proposition-65/proposition-65-faqs.

Radford, Alec, Jeff Wu, Rewon Child, D. Luan, Dario Amodei, and Ilya Sutskever.
    "Language Models Are Unsupervised Multitask Learners," 2019.
    https://www.gwern.net/docs/ai/nn/transformer/gpt/2019-radford.pdf.

"Regulation (EU) 2016/679 of the European Parliament and of the Council of 27
    April 2016 on the Protection of Natural Persons with Regard to the Processing
    of Personal Data and on the Free Movement of Such Data, and Repealing
    Directive 95/46/EC (General Data Protection Regulation)," n.d.
    https://eur-lex.europa.eu/eli/reg/2016/679/oj.

Relihan, Tom. "Will Regulating Big Tech Stifle Innovation?" MIT Sloan, September
    27, 2018.
    https://mitsloan.mit.edu/ideas-made-to-matter/will-regulating-big-tech-stifle-
    innovation.

Rio, Giulia McDonnell Nieto del. "University of California Will No Longer Consider
    SAT and ACT Scores." *The New York Times*, May 15, 2021, sec. U.S.

https://www.nytimes.com/2021/05/15/us/SAT-scores-uc-university-of-california.html.

Salas, Rudy. Artificial intelligence 2019 AB-594 (n.d.).
https://leginfo.legislature.ca.gov/faces/billHistoryClient.xhtml?bill_id=201920200AB594.

Sevilla, Jaime, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos. "Compute Trends Across Three Eras of Machine Learning," 2022. https://doi.org/10.48550/ARXIV.2202.05924.

Shum, Harry. "Advancing Our Ambition to Democratize Artificial Intelligence." The Official Microsoft Blog, November 15, 2016.
https://blogs.microsoft.com/blog/2016/11/15/advancing-ambition-democratize-artificial-intelligence/.

Siebeneck, Todd, and Catherine Wang. "Gross Domestic Product by State, 1st Quarter 2022." Bureau of Economic Analysis. United States Department of Commerce, June 30, 2022.
https://www.bea.gov/sites/default/files/2022-06/qgdpstate0622.pdf.

Singh, Navdeep. "What You Need to Know about the CCPA and the European Union's GDPR," February 26, 2020.
https://www.americanbar.org/groups/litigation/committees/minority-trial-lawyer/practice/2020/what-you-need-to-know-about-the-ccpa-and-the-european-unions-gdpr/.

Smith, Brad. "Microsoft Will Be Carbon Negative by 2030." The Official Microsoft Blog, January 16, 2020.
https://blogs.microsoft.com/blog/2020/01/16/microsoft-will-be-carbon-negative-by-2030/.

Urban Institute. "State and Local Expenditures." Accessed August 28, 2022.
https://www.urban.org/policy-centers/cross-center-initiatives/state-and-local-finance-initiative/state-and-local-backgrounders/state-and-local-expenditures.

QuantGov. "State RegData - Definitive Edition Dashboard." Accessed August 16, 2022. https://www.quantgov.org/state-regdata-definitive-edition.

Stephens, John. "California Consumer Privacy Act," February 14, 2019.
https://www.americanbar.org/groups/business_law/publications/committee_newsletters/bcl/2019/201902/fa_9/.

Sturmey, S.G. "Ch. 9: Flags of Convenience." In *British Shipping and World Competition*. Liverpool University Press, 2009. https://doi.org/10.5949/liverpool/9780986497322.003.0009.

Sullivan, Paul. "What the Small Player Can Expect When Using a Lobbyist." *The New York Times*, January 25, 2013, sec. Your Money. https://www.nytimes.com/2013/01/26/your-money/what-the-small-player-can-expect-when-using-a-lobbyist.html.

Sustainability (US). "Sustainability in the Cloud." Amazon Web Services. Accessed August 28, 2022. https://sustainability.aboutamazon.com/environment/the-cloud.

The City University of New York. "Testing FAQs." Accessed August 28, 2022. https://www.cuny.edu/academics/testing/testing-faqs/.

Congressional Budget Office. "The Federal Budget in 2019: An Infographic," April 15, 2020. https://www.cbo.gov/publication/56324.

Thoppilan, Romal, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, et al. "LaMDA: Language Models for Dialog Applications," 2022. https://doi.org/10.48550/ARXIV.2201.08239.

Google Sustainability. "Tracking Our Carbon-Free Energy Progress." Accessed August 20, 2022. https://sustainability.google/progress/energy/.

Trébaol, Tristan, ed. *CUMULATOR — a Tool to Quantify and Report the Carbon Footprint of Machine Learning Computations and Communication in Academia and Healthcare*, 2020.

UN News. "Urgent Action Needed over Artificial Intelligence Risks to Human Rights," September 15, 2021. https://news.un.org/en/story/2021/09/1099972.

Statista. "U.S. Fortune 500 Companies 2021, by State," September 8, 2021. https://www.statista.com/statistics/303696/us-fortune-500-companies-by-state/.

Vogel, David. *Trading up: Consumer and Environmental Regulation in a Global Economy*. Cambridge, Mass: Harvard University Press, 1995.

Williams, Greg. "DeepMind Is Opening a Huge New London Headquarters in 2020." *Wired UK*. Accessed August 25, 2022. https://www.wired.co.uk/article/deepmind-new-london-headquarters.

ChatBot. "Your Ultimate Chatbot Best Practices Guide." Accessed August 20, 2022. https://www.chatbot.com/chatbot-best-practices/.

Yuan, Binhang, Yongjun He, Jared Quincy Davis, Tianyi Zhang, Tri Dao, Beidi
    Chen, Percy Liang, Christopher Re, and Ce Zhang. "Decentralized Training of
    Foundation Models in Heterogeneous Environments." arXiv, June 10, 2022.
    http://arxiv.org/abs/2206.01288.

Zhang, Baobao, and Allan Dafoe. "Artificial Intelligence: American Attitudes and
    Trends." SSRN Scholarly Paper. Rochester, NY, January 9, 2019.
    https://doi.org/10.2139/ssrn.3312874.

Zhang, Daniel, Nestor Maslej, Erik Brynjolfsson, John Etchemendy, Terah Lyons,
    James Manyika, Helen Ngo, Juan Carlos Niebles, Michael Sellitto, Ellie
    Sakhaee, Yoav Shoham, Jack Clark, and Raymond Perrault, "The AI Index 2022
    Annual Report," AI Index Steering Committee, Stanford Institute for
    Human-Centered AI, Stanford University, March 2022.